Contents

1 Integrability of distributions

Definition 1.1. Let M^n be a smooth manifold of dimension n. A *k-dimensional foliation* on M is a smooth atlas $\varphi_U : U \subseteq \mathbb{R}^k \times \mathbb{R}^{n-k} \to M$, so that all transition maps $\psi_{UV} = \varphi_V^{-1} \circ \varphi_U|_{\varphi_U^{-1}(\varphi_V(V))}$ are of the form $\psi_{UV}(x, y) = (\psi_1(x, y), \psi_2(y))$, where $(x, y) \in \mathbb{R}^k \times \mathbb{R}^{n-k}$.

Notice that the transition maps of a foliation $\mathcal{F}$ preserve the k-dimensional submanifolds $\{y = \text{constant}\}$, and therefore a foliation defines a decomposition of M as a (setwise disjoint, but not topologically disjoint) union of k-dimensional submanifolds, called the *leaves* of $\mathcal{F}$. Given a foliation $\mathcal{F}$ on M a diffeomorphism $f : M \to N$ defines a foliation $f_*(\mathcal{F})$ on N in the obvious way. We say that two foliations $\mathcal{F}_1$ and $\mathcal{F}_2$ on M_1 and M_2 respectively are diffeomorphic if there is a diffeomorphism $f : M_1 \to M_2$ so that $f_*(\mathcal{F}_1) = \mathcal{F}_2$. It's easy to check that any diffeomorphism $f : M_1 \to M_2$ which sends the leaves of $\mathcal{F}_1$ to those of $\mathcal{F}_2$ is a diffeomorphism of foliations.

Example 1.2. If $M = M_1^k \times M_2^l$, there is clearly an l-dimensional foliation $\mathcal{F}_1$ on M whose leaves are $\{x = \text{constant} \in M_1\}$. Similarly there is a k-dimensional foliation $\mathcal{F}_2$ whose leaves are $\{y = \text{constant} \in M_2\}$.

More generally, let $f : M \to N$ be a fiber bundle, ie a proper submersion. Then there is a foliation $\mathcal{F}$ whose leaves are $f^{-1}(\text{constant})$. The dimension of the foliation is $\dim M - \dim N$.

On the torus $T^n = S^1 \times \ldots \times S^1$, consider the 1-dimensional foliation $\mathcal{F}_{(a_1,\ldots,a_n)}$ whose leaves are given by $(\theta_1, \ldots \theta_n) = t(a_1, \ldots, a_n) + \overrightarrow{\text{constant}}$ for some fixed numbers $a_1, \ldots, a_n$. Unless all of the a_j are rational multiples of each other, the leaves will be diffeomorphic to $\mathbb{R}$ rather than S^1. *Leaves of a foliation will typically not be properly embedded.* If all of the a_j are mutually irrational, then any single leaf will be dense in T^n. What is the closure of a leaf if some of the a_j are rational multiples and others are not?

Definition 1.3. Let M^n be a smooth manifold. A *k-dimensional distribution* ξ^k is a smooth choice of k-dimensional plane in TM_p for all points $p \in M$.

Note. To make the definition more precise, note that for any m-vector bundle $E \to M$ we can form the *bundle of k-planes in E*, denoted by $\mathrm{Gr}_k(E) \to M$. This is a fiber bundle whose fiber is $\mathrm{Gr}_{m,k} = \{k - \text{planes in } \mathbb{R}^m\} = \mathrm{GL}_m(\mathbb{R})/(\mathrm{GL}_k(\mathbb{R}) \times \mathrm{GL}_{m-k}(\mathbb{R}))$. Then a k-dimensional distribution is simply a smooth section $\sigma : M \to \mathrm{Gr}_k(TM)$.

Given a foliation $\mathcal{F}$ the distribution $T\mathcal{F}$ is defined by the tangent fibers of the leaves. In fact this distribution determines $\mathcal{F}$ uniquely:

Proposition 1.4. *Let $\mathcal{F}_1, \mathcal{F}_2$ be foliations of equal dimension on M, and suppose $f : M \to M$ is a diffeomorphism so that $f_*(T\mathcal{F}_1) = T\mathcal{F}_2$. Then f is a diffeomorphism of foliations from $\mathcal{F}_1$ to $\mathcal{F}_2$.*

Proof: If f maps local pieces of leaves to local pieces of leaves everywhere, then clearly f maps leaves to leaves globally. Given $p \in M$, choose open sets U

containing p and V containing $f(p)$ which come from the atlases defining $\mathcal{F}_1$ and $\mathcal{F}_2$ respectively. The leaves near p are given by $\{y = \text{constant}\} \subseteq U$ and the leaves near $f(p)$ are given by $\{y = \text{constant}\} \subseteq V$. Therefore at each point $T\mathcal{F}_1 = \text{span}\{\partial_{x_j}\}$ and similarly for $T\mathcal{F}_2$. In coordinates $f_*(T\mathcal{F}_1) = T\mathcal{F}_2$ means that $\frac{\partial f}{\partial x_j}$ is a vector in $\text{span}\{\partial_{x_j}\}$. It follows that $f(x,y) = (f_1(x,y), f_2(y))$, that is, f maps $\mathcal{F}_1$ to $\mathcal{F}_2$. □

Note. The coordinates in the above proof may perhaps obscure the central idea, which intuitively is this. The leaf of any foliation locally looks like a graph of a function $g : \mathbb{R}^k \to \mathbb{R}^{n-k}$. For f to preserve $T\mathcal{F}$ means exactly that the first derivatives of g are preserved. But any two functions passing through the same point with equal first derivatives are equal.

Definition 1.5. Let ξ be a distribution on a manifold M. We say that ξ is *integrable* if $\xi = T\mathcal{F}$ for some foliation $\mathcal{F}$ on M (which by the previous proposition is uniquely defined by ξ if it exists).

Proposition 1.6. *Let ξ be an integrable distribution on M. Then for any vector fields X, Y which are tangent to ξ, $[X,Y]$ is tangent to ξ as well.*

Proof: The statement is local. Since ξ is integrable it follows that in foliated coordinates $\xi = \text{span}\{\partial_{x_j}\}$, and therefore $X = g^j(x,y)\partial_{x_j}$ and $Y = h^j(x,y)\partial_{x_j}$ (standard summation convention). We simply calculate $[X,Y] = XY - YX = g^j \frac{\partial h^i}{\partial x_j}\partial_{x_i} - h^j \frac{\partial g^i}{\partial x_j}\partial_{x_i} \in \xi$. □

Note. The word integrability means exactly what it means in vector calculus. As described in the previous note, foliations $\mathcal{F}$ are locally given as graphs of functions, and $T\mathcal{F}$ is the data of the first derivatives of those functions. Asking then if a given distribution is integrable is locally the same as asking for functions with given first derivatives. The previous proposition describes an obvious necessary condition: second order mixed partial derivatives must be equal.

The converse of the above proposition is much more interesting, it says that the only obstruction to integrability is first order and local.

Theorem 1.7 (Frobenius Integrability)**.** *Let ξ be a distribution on M with the property that $[X,Y] \in \Gamma(\xi)$ for all $X, Y \in \Gamma(\xi)$. Then ξ is integrable. (We use the notation $\Gamma(\xi) \subseteq \Gamma(TM)$ for all vector fields tangent to ξ.)*

We first prove a

Lemma 1.8. *Given a vector field X, let $\varphi_t^X : M \to M$ be the flow of X (perhaps only defined for small t and subsets of M if X is not complete). Given $X, Y \in \Gamma(TM)$, $[X,Y] = 0$ if and only if $\varphi_t^X \circ \varphi_s^Y = \varphi_s^Y \circ \varphi_t^X$.*

Proof: Notice that for any diffeomorphism f, $\varphi_t^{f_*(X)} = f \circ \varphi_t^X \circ f^{-1}$ (by the chain rule, and uniqueness of flows). Our lemma reduces then to $[X,Y] = 0$ iff $X = \varphi_{s*}^Y(X)$ for all s. Differentiating $\varphi_{s*}^Y(X) - X$ with respect to s at $s = 0$ gives the definition of $\mathcal{L}_Y X$, which is equal to $-[X,Y]$. Since $\varphi_{s_0+s}^Y =$

$\varphi^Y_{s_0} \circ \varphi^Y_s$, if we instead differentiate at $s = s_0$ we get $\varphi^Y_{s_0 *}[X, Y]$, but since $\varphi^Y_{s_0}$ is a diffeomorphism $\varphi^Y_{s_0 *}$ sends only the vector 0 to 0. □

Proof of Frobenius Integrability: First, notice that Proposition 1.4 implies that integrability is a priori a local question: if ξ is integrable on two open sets which intersect, then the foliation must agree on the overlap and it follows that we can glue.

For a point $p \in M$ we therefore choose coordinates so that $\xi_p = \text{span}\{\partial_{y_j}\}_{j=1}^k$ (at the point $p = 0 \in \mathbb{R}^n$ only!) (since we have no foliation yet we denote coordinates on $\mathbb{R}^n$ by $(y_1, \ldots, y_n)$). Let $\pi : \mathbb{R}^n \to \mathbb{R}^k$ be the orthogonal projection. Then $\pi_*|_{\xi_p} : \xi_p \to T\mathbb{R}^k$ is an isomorphism (in fact it is the identity), it follows by continuity of ξ that $\pi_*|_{\xi_q} : \xi_q \to T\mathbb{R}^k$ will be an isomorphism for all nearby q (and thus by shrinking our chart we assume it is true everywhere). Define $X^j_q = \pi_*|^{-1}_{\xi_q}(\partial_{y_j})$. This defines smooth vector fields X^j on our chart, and furthermore $\pi_*[X^j, X^i] = [\pi_* X^j, \pi_* X^i] = [\partial_{y_j}, \partial_{y_i}] = 0$. By assumption $[X^j, X^i]_q \in \xi_q$, and since π_* is an isomorphism on this subspace it follows that $[X^j, X^i] = 0$ everywhere.

We now define new coordinates on our chart by $(x_1, \ldots, x_k, y_{k+1}, \ldots, y_n) \mapsto \varphi^{X^1}_{x_1} \circ \ldots \circ \varphi^{X^k}_{x_k}(0, \ldots, 0, y_{k+1}, \ldots, y_n)$. We see immediately that $\partial_{x_1} = X^1$. But since $[X^j, X^i] = 0$ the lemma implies than all of these flows commute, and therefore $\partial_{x_j} = X^j$. That is to say, $\text{span}\{\partial_{x_j}\} = \xi$, which is to say that these coordinates define a foliation $\mathcal{F}$ with $T\mathcal{F} = \xi$. □

Definition 1.9. Let ξ be a distribution on M. Define $\omega_\xi : \Gamma(\xi) \otimes \Gamma(\xi) \to \Gamma(TM/\xi)$ to simply be the induced operation from $[\cdot, \cdot]$, that is, $\omega_\xi = \pi \circ [\cdot, \cdot]|_\xi$, where $\pi : TM \to TM/\xi$ is the projection.

This definition allows us to state the Frobenius theorem very succinctly: a distribution ξ is integrable if and only if $\omega_\xi \equiv 0$. An important property of ω_ξ is that it is tensorial:

Proposition 1.10. *ω_ξ is $C^\infty(M)$-linear.*

For $f \in C^\infty(M)$, $X, Y \in \Gamma(\xi)$, $\omega_\xi(X, fY) = \pi[X, fY] = \pi(f[X, Y] + X(f)Y) = f\pi[X, Y] + X(f)\pi(Y) = f\omega_\xi(X, Y) + 0$. Since $\omega_\xi(X, Y) = -\omega_\xi(Y, X)$ (a property inherited from $[\cdot, \cdot]$) it is also C^∞ linear in the first slot. □

This implies that for $p \in M$, $\omega_\xi(X, Y)_p$ only depends on X_p and Y_p, that is, ω_ξ is a smooth family of linear maps $\omega_{\xi,p} : \xi_p \otimes \xi_p \to (TM/\xi)_p$. The advantage of this is that it turns integrability into an infinitesimal question rather than a local question: it makes sense to say that ξ is integrable at a point p (which is NOT the same thing as ξ being integrable in a neighborhood of p). It also gives us a way to measure “how non-integrable” a distribution is: because $\omega_{\xi,p}$ is a map of finite dimensional vector spaces, we can think of ξ as being “nearly integrable” if the rank of $\omega_{\xi,p}$ is small, or “very non-integrable” if the rank is large.

Definition 1.11. Let ξ^{n-1} be a hyperplane distribution on a smooth manifold M^n. We say that ξ is a *contact structure* if $\omega_{\xi p}$ is non-degenerate (as a bilinear form) for every point $p \in M$.

Often times this definition is stated simply as "a contact structure is a maximally non-integrable hyperplane field". To be more concrete, suppose we choose a metric on M. This gives a canonical (up to $\pm$) identification of $(TM/\xi)_p \cong \mathbb{R}$, and then $\omega_{\xi,p}$ can be presented as $\omega_{\xi,p}(X_p, Y_p) = \langle A_p X_p, Y_p \rangle$ for some linear map $A_p : \xi_p \to \xi_p$. To say that $\omega_{\xi,p}$ is non-degenerate simply means that the determinant of A_p is non-zero. Notice also that A_p is skew-symmetric. Before we explore the geometry of contact structures it's essential to have a good understanding of the linear structure.

2 Symplectic linear algebra

Definition 2.1. A *symplectic structure* ω on a (finite dimensional) vector space V is a skew, non-degenerate bilinear map.

Proposition 2.2. *Let ω be a skew bilinear map on a vector space V. The following are equivalent characterizations of non-degeneracy:*

- *for every nonzero $v \in V$ there is $w \in V$ so that $\omega(v, w) \neq 0$*
- *the map $\omega^* : V \to V^*$ defined by $\omega^*(v) = \omega(v, \cdot)$ is an isomorphism*
- *$\omega \wedge \ldots \wedge \omega \neq 0$, where the number of terms on the left hand side is $\frac{1}{2} \dim V$.*

In particular, the third condition implies that any vector space with a symplectic structure is even dimensional.

Proof: The first and second conditions are obviously equivalent. To show that the third condition implies the first, suppose we have a vector $v \in V$ so that $\omega(v, w) = 0$ for all $w \in V$. Let $\{v, v_2, \ldots, v_{2n}\}$ be any basis for V including the vector v. Then by definition, $\omega^{\wedge n}(v, v_2, \ldots, v_{2n}) = \frac{1}{(2n)!} \sum_{\sigma \in S_{2n}} \mathrm{sign}(\sigma) \omega(v_{\sigma(1)}, v_{\sigma(2)}) \ldots \omega(v_{\sigma(2n-1)}, v_{\sigma(2n)})$, where S_{2n} is the group of permutations on $2n$ elements. We see that every term in the sum is zero, and since $\omega^{\wedge n}$ is a top dimensional alternating tensor it is determined by its value on any basis.

Finally we show that the first condition implies the third; to do this, we construct a basis $\{v_1, \ldots, v_{2n}\}$ of V so that $\omega^{\wedge n}$ is nonzero on this basis. Choose $v_1 \in V$ arbitrarily, and then choose v_2 so that $\omega(v_1, v_2) \neq 0$. We then define the subspace $V_2 \subseteq V$ to be $V_2 = \{v \in V; \omega(v, v_1) = \omega(v, v_2) = 0\}$. First note that $\dim V_2 = \dim V - 2$: by definition $V_2 = \ker \omega^*(v_1) \cap \ker \omega^*(v_2)$, and since ω^* is an isomorphism we know the covectors $\omega^*(v_1)$ and $\omega^*(v_2)$ are linearly independent. We further claim that $\omega|_{V_2}$ also satisfies the first property of the proposition. Indeed given $v \in V_2$, we can choose $w \in V$ so that $\omega(v, w) \neq 0$. But v_1 and v_2 span V/V_2, and therefore $w = av_1 + bv_2 + w'$ for some real numbers a, b and $w' \in V_2$; it follows that $\omega(v, w') \neq 0$.

We then choose $v_3 \in V_2$ arbitrarily, $v_4 \in V_2$ satisfying $\omega(v_3, v_4) \neq 0$, define $V_3 = \{v \in V_2; \omega(v, v_3) = \omega(v, v_4) = 0\}$, and continue the process inductively. Because $\dim V_{k+1} = \dim V_k - 2$, we must either end up with a V_n whose dimension is either 1 or 0. But $\dim V_n = 1$ is impossible: we proved that $\omega|_{V_n}$ satisfies the first property of the proposition, but every skew bilinear form on a 1-dimensional vector space obviously vanishes. Therefore V_n is zero dimensional and thus this process defines a basis $\{v_1, \dots, v_{2n}\}$. By construction, $\omega(v_i, v_j) = 0$ unless $\{i, j\} = \{2k-1, 2k\}$ for some k. Then $\omega^{\wedge n}(v, v_2, \dots, v_{2n}) = \frac{2^n}{(2n)!}\omega(v_1, v_2) \dots \omega(v_{2n-1}, v_{2n}) \neq 0$. □

Definition 2.3. Let V be a vector space with a symplectic structure ω. Then a subspace $L \subseteq V$ is called *isotropic* if $\omega|_L = 0$.

Proposition/Definition 2.4. Let L be an isotropic subspace in (V, ω). Then $\dim L \leqslant \frac{1}{2} \dim V$. If $\dim L = \frac{1}{2} \dim V$ we say that L is *Lagrangian*.

Proof: Let $\{v_1, \dots, v_k\}$ be a basis for L, and let $\eta_j = \omega^*(v_j)$. Let $C = \ker \eta_1 \cap \dots \cap \ker \eta_k$. Because ω^* is an isomorphism $\{\eta_j\}$ is a linearly independent set in V^*, therefore $\dim C = \dim V - k$. Notice that the isotropic condition is equivalent to $L \subseteq C$, which implies $k \leqslant \dim V - k$. □

Example 2.5. Let $V = \mathbb{R}^{2n}$ with coordinates $(x_1, y_1, \dots, x_n, y_n)$, and let $\omega_{\text{std}} = \sum_{i=1}^n x_i^* \wedge y_i^*$. Then $\omega^{\wedge n} = x_1^* \wedge y_1^* \wedge \dots \wedge x_n^* \wedge y_n^* = \det \neq 0$ and so ω is a symplectic structure. The subspace $L_0 = \{x_j = 0\}$ is a Lagrangian. If $A : \mathbb{R}^n \to \mathbb{R}^n$ is any symmetric matrix then $L = \{y = Ax\}$ is also a Lagrangian. Check that every Lagrangian which is transverse to L_0 is of this form (the graph of a symmetric matrix). Also check that this example is essentially unique: every symplectic vector space (V, ω) is isomorphic to $(\mathbb{R}^{2n}, \omega_{\text{std}})$.

3 Contact structures defined by differential forms

For many purposes (especially computational), it is often useful to write a hyperplane distribution ξ as the kernel of a non-zero 1-form $\alpha \in \Omega^1(M) := \Gamma(TM^*)$. This can always be done locally. Globally there is an orientability obstruction: the bundle TM/ξ is a 1-dimensional vector bundle, and α descends to a non-vanishing section of $(TM/\xi)^*$, implying that this bundle (and its dual) is trivial. Conversely, if ξ is a hyperplane distribution so that TM/ξ is orientable, we can choose a non-vanishing section $V \in \Gamma(TM/\xi)$ (since any orientable 1-dimensional vector bundle is trivial). We can then $\alpha \in \Omega^1(M)$ by $\alpha(V) = 1$ and $\alpha|_\xi = 0$ so that $\xi = \ker \alpha$. We say that a hyperplane distribution is *coorientable* if this property holds.

If ξ is a contact structure on M, call $\alpha \in \Omega^1(M)$ a *contact form* for ξ if $\ker \alpha = \xi$. The main goal of this section is the following

Proposition 3.1. *Let $\alpha \in \Omega^1(M) = \Gamma(TM^*)$ be a nowhere vanishing 1-form on a smooth $(2n+1)$-manifold M. Then $\xi := \ker \alpha$ is a contact structure if and only if $\alpha \wedge (d\alpha)^{\wedge n}$ is a volume form (ie a nowhere vanishing element of $\Omega^{2n+1}(M)$).*

Note. When $2n + 1 = 4k + 3$ the sign of $\alpha \wedge (d\alpha)^{\wedge n}$ does not depend on the sign of α. Thus, *a contact structure on M^{4k+3} (coorientable or not) defines a canonical orientation on M.* When $\dim M = 4k + 1$ a coorientation of ξ is necessary to define an orientation on M, and indeed non-coorientable contact structures can exist only on non-orientable manifolds.

There are two main ingredients in Proposition 3.1, one is Proposition 2.2, and the other is

Proposition 3.2 (Cartan's formula)**.** *Let $\eta \in \Omega^k(M)$ be a differential form, and let $X \in \Gamma(TM)$ be a vector field. Then $\mathcal{L}_X\eta = X \lrcorner\, d\eta + d(X \lrcorner\, \eta)$.*

Proof: First, notice that

$$X \lrcorner\, d(\eta_1 \wedge \eta_2) + d(X \lrcorner\, (\eta_1 \wedge \eta_2))$$

$$= X \lrcorner\, (d\eta_1 \wedge \eta_2 + (-1)^k \eta_1 \wedge d\eta_2) + d(X \lrcorner\, \eta_1 \wedge \eta_2 + (-1)^k \eta_1 \wedge X \lrcorner\, \eta_2)$$

$$= X \lrcorner\, d\eta_1 \wedge \eta_2 + (-1)^{k+1} d\eta_1 \wedge X \lrcorner\, \eta_2 + (-1)^k X \lrcorner\, \eta_1 \wedge d\eta_2 + (-1)^k(-1)^k \eta_1 \wedge X \lrcorner\, d\eta_2$$

$$+ d(X \lrcorner\, \eta_1) \wedge \eta_2 + (-1)^{k-1} X \lrcorner\, \eta_1 \wedge d\eta_2 + (-1)^k d\eta_1 \wedge X \lrcorner\, \eta_2 + (-1)^k(-1)^k \eta_1 \wedge d(X \lrcorner\, \eta_2)$$

$$= (X \lrcorner\, d\eta_1 + d(X \lrcorner\, \eta_1)) \wedge \eta_2 + 0 + 0 + \eta_1 \wedge (X \lrcorner\, d\eta_2 + d(X \lrcorner\, \eta_2)).$$

Since both operators $\mathcal{L}_X$ and $X \lrcorner\, d + dX \lrcorner$ are Leibnizian with respect to $\wedge$ it suffices to demonstrate their equality for 1-forms η. Also since it is local it suffices to work in a chart. Suppose $X \neq 0$, and choose coordinates where $X = \partial_{x_1}$ and write $\eta = f^j dx_j$ for some smooth functions f^j. $\mathcal{L}_X\eta = \frac{\partial f^j}{\partial x_1} dx_j$ and $d(X \lrcorner\, \eta) = df^1 = \frac{\partial f^1}{\partial x_j} dx_j$. $X \lrcorner\, d\eta = \partial_{x_1} \lrcorner\, \frac{\partial f^j}{\partial x_i} dx_i \wedge dx_j = \frac{\partial f^j}{\partial x_1} dx_j - \frac{\partial f^1}{\partial x_i} dx_i$.

This establishes Cartan's formula for non-vanishing vector fields. Finally we note that, since the formula is linear in X and any vector field can be locally written as the sum of non-vanishing vector fields, the result follows. □

Proof of Proposition 3.1: By choosing a (local) basis for TM so that all but one element is contained in $\xi = \ker \alpha$, we see that $\alpha \wedge (d\alpha)^{\wedge n}$ is nonvanishing if and only if $(d\alpha)^{\wedge n}|_\xi$ is nonvanishing. Applying Proposition 2.2 we see this is equivalent to $d\alpha|_\xi$ being non-degenerate at each point. We claim that $d\alpha(X,Y) = -\alpha(\omega_\xi(X,Y))$ for any vector fields X, Y tangent to ξ. This claim completes the proof since then ω_ξ is non-degenerate exactly when $d\alpha|_\xi$ is.

To prove the claim, recall that the Lie derivative is Leibnizian with respect to evaluation of 1-forms on vector fields: $\mathcal{L}_X(\alpha(Y)) = (\mathcal{L}_X\alpha)(Y) + \alpha(\mathcal{L}_X Y)$. In our case the left hand side is zero (since $\alpha(Y) = 0$), the right hand side equals $(X \lrcorner\, d\alpha + d(\alpha(X)))\,(Y) + \alpha([X,Y]) = (X \lrcorner\, d\alpha + d(0))(Y) + \alpha([X,Y]) = d\alpha(X,Y) + \alpha([X,Y])$. □

Example 3.3. On $\mathbb{R}^{2n+1}$ with coordinates $(x_1, y_1, \ldots, x_n, y_n, z)$, let $\alpha_{\mathrm{std}} = dz - \sum_{i=1}^n y_i dx_i$. Then $\xi_{\mathrm{std}} = \ker \alpha_{\mathrm{std}}$ is contact, since $\alpha \wedge (d\alpha)^{\wedge n}$ is the standard volume form.

Alternatively, let $\alpha_{\mathrm{rot}} = dz - \sum_{i=1}^n y_i dx_i - x_i dy_i$. This contact form is rotationally symmetric: with respect to polar coordinates $(x_i, y_i) = (r_i \cos\theta_i, r_i \sin\theta_i)$, we have $\alpha_{\mathrm{rot}} = dz + \sum_{i=1}^n r_i^2 d\theta_i$. In fact this example is *contactomorphic* to the previous one: there exists a diffeomorphism $f : \mathbb{R}^{2n+1} \to \mathbb{R}^{2n+1}$ so that $f(\xi_{\mathrm{std}}) = \xi_{\mathrm{rot}}$.

On T^3, let $\alpha_1 = \cos(y)dz - \sin(y)dx$ $(x, y, z \in \mathbb{R}/2\pi\mathbb{Z})$. Then $\xi_1 = \ker\alpha_1$ is contact. More generally, let $\alpha_k = \cos(ky)dz - \sin(ky)dx$ for any $k \in \mathbb{Z}_{>0}$. These all define contact structures, and are mutually non-contactomorphic. Notice furthermore that if we define $\alpha_{1/2} = \cos(\frac{y}{2})dz - \sin(\frac{y}{2})dx$, while this is not a well-defined 1-form, it does have a well defined kernel. $\xi = \ker\alpha_{1/2}$ is an example of a non-coorientable contact structure.

4 Contact isotopies

4.1 Contact vector fields

At first, the definition of a contact structure may seem quite strange. Like any geometry, the best way to get an intuitive understanding of it is to study its symmetries.

Definition 4.1. Let ξ_1, ξ_2 be contact structures on M_1 and M_2 respectively. A diffeomorphism $\varphi : M_1 \to M_2$ is called a *contactomorphism* if $\varphi_*(\xi_1) = \xi_2$. A *contact isotopy* is an isotopy through contactomorphisms, that is, a smooth 1-parameter family of contactomorphisms $\varphi_t : M_1 \to M_1$ so that $\varphi_0 = \mathrm{id}$ and $\varphi_{t*}(\xi_1) = \xi_1$.

The nice thing about contact isotopies rather than contactomorphisms is that we can view them as flows of vector fields.

Definition 4.2. Let (M, ξ) be a contact manifold. We say $X \in \Gamma(TM)$ is a *contact vector field* if $[X, Y] \in \Gamma(\xi)$ for all vector fields $Y \in \Gamma(\xi)$.

Remark 4.3. A contact vector field is not the same thing as a vector field tangent to ξ. We call this latter type of vector field *tangent* vector fields. Thus the defining property of contact vector fields X is that $\mathcal{L}_X$ sends tangent vector fields to tangent vector fields. In fact, these two types of vector field are somewhat opposites: **a contact vector field is never tangent to ξ on an open set**. Indeed, the non-degeneracy property defining a contact structure says exactly that.

Proposition 4.4. *Let φ_t be a contact isotopy on a contact manifold (M, ξ). Then $X^t := \frac{d\varphi_t}{dt} \circ \varphi_t^{-1}$ is a contact vector field for all $t \in [0, 1]$. Conversely, let X^t be a smooth family of contact vector fields. Then the flow of X^t is a contact isotopy.*

Proof: Suppose φ_t is a contact isotopy, and let Y be a vector field tangent to ξ. Then $\varphi_{t*}(Y)_p \in \xi_p$ at each point $p \in M$, by definition of contact isotopy. So $\varphi_{t+h*}(Y)_p - \varphi_{t*}(Y)_p \in \xi_p$ as well, and but putting $\lim\limits_{h\to 0} \frac{1}{h}$ in front of this equation we get $\mathcal{L}_{\frac{d\varphi_t}{dt}}(Y \circ \varphi_t) = [\frac{d\varphi_t}{dt}, Y \circ \varphi_t] = [X^t, Y] \circ \varphi_t \in \xi$.

Now suppose X^t is given, then the flow φ_{X^t} is defined by the equation $\frac{d\varphi_{X^t}}{dt} \circ \varphi_t^{-1} = X^t$ (together with $\varphi_{X^0} = \mathrm{id}$). For a given point $p \in M$ and a vector field $Y \in \Gamma(\xi)$, let $c : [0, 1] \to (TM/\xi)_p$ be defined by $c(t) = \pi\varphi_{X^t*}(Y)_p$ (where

$\pi : TM \to TM/\xi$ is the projection). Then $c(0) = 0$ and $c'(t) = \pi([X^t, Y]) = 0$. Thus $c(t) = 0$ for all t which implies that φ_{X^t} is a contact isotopy. □

Because of this proposition, we use the notation $T\,\mathrm{Cont}(M,\xi)_{\mathrm{id}}$ to denote the space of all contact vector fields (read literally as the tangent fiber to the group of contactomorphisms). This is clearly a vector space, though be careful to note that a smooth function times a contact vector field is not typically a contact vector field.

Theorem 4.5. *Let (M,ξ) be a contact manifold. Let $\pi : T\,\mathrm{Cont}(M,\xi)_{\mathrm{id}} \to \Gamma(TM/\xi)$ be the restriction of the projection $\Gamma(TM) \to \Gamma(TM/\xi)$. Then π is an isomorphism of vector spaces.*

Proof: π is linear so to show that it is injective we show it has trivial kernel. If $\pi(X) = 0$ then X is tangent to ξ everywhere. If $X \neq 0$ then the contact condition implies there is a vector field Y tangent to ξ so that $[X,Y] \notin \xi$, which contradicts the fact that X is contact.

To show π is surjective, suppose we are given a section $\sigma : M \to \Gamma(TM/\xi)$. Choose a random (not contact) vector field Z so that $\pi(Z) = \sigma$. Define a function $E : \Gamma(\xi) \to \Gamma(TM/\xi)$ by $E(Y) = \pi([Z,Y])$, notice that Z is a contact vector field iff $E = 0$. It's easy to see that E is $C^\infty(M)$-linear (just as in the proof of Proposition 1.10), so in fact E descends to a map $E : \xi \to TM/\xi$.

The contact condition (interpreted as the second property in Proposition 2.2) implies that there is a vector field $V = (\omega_\xi^*)^{-1}(E)$ so that V is tangent to ξ and $\omega_\xi(V,\cdot) = E$. We let $X = Z - V$. Then for any $Y \in \Gamma(\xi)$, $\pi([X,Y]) = \pi([Z,Y]) - \pi([V,Y]) = E(Y) - E(Y) = 0$. Therefore X is contact and $\pi(X) = \sigma$ (since $V \in \Gamma(\xi)$), so π is surjective. □

The power of this theorem is that it shows us that contact isotopies are easy to construct: the space of contact isotopies is infinite dimensional, and admits cutoff functions (ie they are locally constructable). For example:

Corollary 4.6 (isotopy extension)**.** *Let (M,ξ) and (U,ζ) be two contact manifolds of equal dimension. Let $f_t : U \to M$ be a smooth path of contact embeddings. Then there is a contact isotopy φ_t on M, so that $\varphi_t \circ f_0 = f_t$. Furthermore, we can assume that φ_t is the identity outside any open set $V_t \subseteq M$ containing the closure of $f_t(U)$.*

Proof: $X_p = \frac{df_t(p)}{dt}$ defines a contact vector field on M which is defined over $f_t(U)$. Using Theorem 4.5 we can represent this vector field as a section $\sigma_t : f_t(U) \to (TM/\xi)|_{f_t(U)}$. We can then extend this section to all of M by cutting it off to zero (in a neighborhood of V_t). The flow of this section defines the desired φ_t. □

4.2 Contact isotopy conditions with a contact form

We would also like to understand contact isotopies and vector fields from the point of view of contact forms.

Proposition 4.7. *Let $(M, \ker\alpha)$ be a contact manifold with contact form α. A diffeomorphism $\varphi : M \to M$ is a contactomorphism if and only if $\varphi^*\alpha = f\alpha$ for some nowhere vanishing function $f \in C^\infty(M)$. A vector field X is a contact vector field if and only if $\mathcal{L}_X\alpha = g\alpha$ for some (possibly zero) function $g \in C^\infty(M)$.*

Proof: The first statement is obvious: by definition of pullback we have that $(\varphi^*\alpha)(Y) = \alpha(\varphi_*Y)$. For the second statement, let Y be a vector field tangent to ξ. Then $0 = \mathcal{L}_X(\alpha(Y)) = (\mathcal{L}_X\alpha)(Y) + \alpha([X,Y])$. X being contact means that $\alpha([X,Y]) = 0$ for all $Y \in \ker\alpha$, therefore it follows that $\ker\mathcal{L}_X\alpha \supseteq \ker\alpha$. □

Proposition/Definition 4.8. Let $(M, \ker\alpha)$ be a contact manifold with contact form α. α induces a map $\alpha : T\operatorname{Cont}(M, \ker\alpha)|_{\mathrm{id}} \to C^\infty M$; this map is an isomorphism of vector spaces. We define the *Reeb vector field* R_α to be the contact vector field that is sent to the constant function $1 \in C^\infty M$. Then $R_\alpha \lrcorner\, d\alpha = 0$.

Proof: The first statement is simply a rephrasing of Theorem 4.5: a 1-form α with $\ker\alpha = \xi$ defines an identification of TM/ξ with $\mathbb{R} \times M$, and under this identification the projection of vector fields $\pi : TM \to TM/\xi$ is the same as evaluation $\alpha : TM \to \mathbb{R}\times M$. R_α is then defined to be the unique contact vector field with $\alpha(R_\alpha) = 1$. Proposition 4.7 says that $\mathcal{L}_{R_\alpha}\alpha$ is a multiple of α. By Cartan's formula this is equal to $R_\alpha \lrcorner\, d\alpha + d(R_\alpha \lrcorner\, \alpha) = R_\alpha \lrcorner\, d\alpha + d(1) = R_\alpha \lrcorner\, d\alpha$. Since $R_\alpha \lrcorner\, d\alpha$ is a multiple of α it is zero when restricted to $\ker\alpha$, but it is also zero on R_α since $d\alpha$ is skew. Thus it is zero on TM. □

4.3 Gray stability

Theorem 4.9 (Gray stability)**.** *Let M be a smooth manifold, and let $\{\xi_t\}_{t\in[0,1]}$ be a homotopy of contact structures on M. Assume that ξ_t is compactly supported ($\xi_t = \xi_0$ for all t outside of a compact set in M). Then there exists a smooth isotopy $\varphi_t : M \to M$ so that $\varphi_{t*}(\xi_0) = \xi_t$.*

Philosophically this theorem should be thought of as saying "contact structures have no deformation theory". Said differently, the space of contact structures on M modulo $\operatorname{Diff}_0(M)$ is discrete.

Corollary 4.10. *Let (M, ξ_0) be a contact manifold, and suppose ξ_1 is another contact structure C^1 close to ξ_0. Then ξ_1 is isotopic to ξ_0.*

Proof: Let ξ_t be a homotopy through distributions between ξ_0 and ξ_1; since ξ_1 is C^1 close to ξ_0 this homotopy can be chosen to be C^1 small. However, the condition on distributions to be contact is C^1 open (ie $\alpha \wedge (d\alpha)^{\wedge n} \neq 0$ is an open condition on α and its first derivative), and therefore ξ_t is contact for all t. Applying Gray stability we get an isotopy $\varphi_t : M \to M$ so that $\varphi_{1*}(\xi_0) = \xi_1$. □

Before proving the theorem, we need to define some basic notions from differential topology: the time derivative of a family of distributions, and the Lie derivative of a distribution.

Definition 4.11. Let ξ be a distribution on M, and let X be a vector field. Define $\mathcal{L}_X\xi : \xi \to TM/\xi$ to be the bundle map defined by $(\mathcal{L}_X\xi)(Y) = \pi(\mathcal{L}_X Y)$ for $Y \in \Gamma(\xi)$. Note that by the standard argument $\mathcal{L}_X\xi$ is tensorial.

Notice that this definition gives us a particularly clean way to define the contact condition: a hyperplane field ξ is contact if and only if every bundle map $F : \xi \to TM/\xi$ can be realized as $F = \mathcal{L}_X\xi$ for some tangent vector field $X \in \Gamma(\xi)$.

Definition 4.12. Let ξ_t be a 1-parameter family of distributions on a manifold M. Suppose that ξ_t is locally given as the span of a basis $\{Y_j^t\}_{j=1,\dots,k}$. Define $\dot{\xi}_t = \frac{d\xi_t}{dt}$ as the linear bundle map $\xi_t \to TM/\xi_t$ which maps $\sum_{j=1}^{k} a_j Y_j^t$ to $\sum_j a_j \pi_t(\dot{Y}_j^t)$, where $\pi_t : TM \to TM/\xi_t$ is the projection. Note that $\dot{\xi}_t$ is well defined independent the choice of $\{Y_j\}$. Also note that ξ_t is constant in t if and only if $\dot{\xi}_t = 0$ for all t.

Proposition 4.13. *Let M be a manifold with a family of distributions ξ_t, and let φ_t be any smooth isotopy, which is a flow of the vector field X^t. Then $\frac{d}{dt}(\varphi_{t*}(\xi_t)) = \varphi_{t*}\left(\mathcal{L}_{X^t}\xi_t + \dot{\xi}_t\right)$.*

Proof: This is an immediate consequence of the corresponding identity for vector fields: $\frac{d}{dt}(\varphi_{t*}(Y^t)) = \varphi_{t*}\left(\mathcal{L}_{X^t}Y^t + \dot{Y}^t\right)$. □

Proof of Gray stability: The idea is to construct the isotopy as the flow of a vector field. Let us define X^t to be the vector field tangent to ξ_t satisfying $\mathcal{L}_{X^t}\xi_t = -\dot{\xi}_t$ which we know exists because of the contact condition. Notice that $X^t = 0$ everywhere ξ_t is constant, so the condition that the homotopy is compactly supported ensures the flow of X^t is well defined; call this flow φ_t. Since $\frac{d}{dt}\varphi_{t*}(\xi_t) = \varphi_{t*}\left(\mathcal{L}_{X^t}\xi_t + \dot{\xi}_t\right) = 0$ we see that $\varphi_{t*}(\xi_t) = \xi_0$ for all t. □

Recall that $\mathbb{R}^{2n+1}$ has a standard contact structure, $\xi_{\text{std}} = \ker\alpha_{\text{std}}$, where $\alpha_{\text{std}} = dz - \sum_i y_i dx_i$. As an application of Gray stability, we show that this model is in fact a universal local model.

Theorem 4.14 (Darboux's theorem)**.** *Let (M,ξ) be a contact manifold. Then given $p \in M$ there is an open neighborhood $U \subseteq M$ containing p and a contact embedding $\varphi : (U, \xi|_U) \to (\mathbb{R}^{2n+1}, \xi_{std})$ so that $\varphi(p) = 0$. Furthermore, given $\xi = \ker\alpha$, we can find a (perhaps smaller) neighborhood $U' \ni p$ and a embedding $\psi : U' \to \mathbb{R}^{2n+1}$ satisfying $\psi^*(\alpha_{std}) = \alpha|_{U'}$ and $\psi(p) = 0$.*

Proof: Since every contact structure locally admits a contact form, the second statement implies the first. First, find a smooth embedding $f : U \to \mathbb{R}^{2n+1}$ so that $f(p) = 0$ and $f^*(\alpha_{\text{std}})_p = \alpha_p$. Define $\alpha^t = (1-t)\alpha + tf^*(\alpha_{\text{std}})$, notice that α^t is constant at the point p. In particular $\ker\alpha^t$ is a contact structure for all

$t \in [0,1]$ at the point p, so by possibly shrinking U we can assume that $\ker \alpha^t$ is a contact structure everywhere on U for all t. We also shrink U so that it has the property that R_{α^t} has no closed orbits anywhere in U.

We would now like to apply Gray's theorem to get an isotopy realizing this homotopy. But our case here is different from the statement of the theorem: we have contact forms rather than contact structures, and non-compactly supported homotopies. So we re-run the argument. Imagine we had an isotopy φ_t so that $\varphi_t^*(\alpha_t) = \alpha_0 = \alpha$. Differentiating this equation gives $\varphi_t^*(\mathcal{L}_{X^t}\alpha_t + \dot{\alpha}_t) = 0$. Dropping the φ_t^* and applying Cartan's formula then gives us $X^t \lrcorner\, d\alpha_t + d(\alpha_t(X^t)) + \dot{\alpha}_t = 0$. We write $X^t = Y^t + h^t R_{\alpha_t}$ where $\alpha_t(Y_t) = 0$. (In our earlier proof of Gray stability, we were able to choose $h^t = 0$ everywhere, but this was not forced on us. It is this additional freedom that we exploit in order to get an isotopy of contact *forms*.)

The above equation becomes $Y^t \lrcorner\, d\alpha_t + dh^t + \dot{\alpha}_t = 0$. Plugging R_{α_t} into this 1-form gives $R_{\alpha_t}(h^t) + \dot{\alpha}_t(R_{\alpha_t}) = 0$. *Because R_{α_t} has no closed orbits* we can always solve this equation for h^t by integration: $h^t = -\int_\gamma \dot{\alpha}_t(R_{\alpha_t})ds$, where $\gamma : (a,b) \to U$ is an orbit of R_{α_t}: that is $\frac{d\gamma}{ds}|_{s=s_0} = R_{\alpha_t,\gamma(s_0)}$. h^t is only defined modulo a function which is constant on all R_{α_t}-orbits: we exploit this freedom to choose $h^t(p) = 0$ and $dh^t_p|_{\ker \alpha_t} = 0$. $dh^t_p(R_{\alpha_t}) = -\dot{\alpha}_t(R_{\alpha_t})_p$ is forced on us, but since $\dot{\alpha}_{tp} = 0$ we in fact have $dh^t_p = 0$ on all of TU_p.

We now return to the equation $Y^t \lrcorner\, d\alpha_t + dh^t + \dot{\alpha}_t = 0$ having solved for h^t. Restricting to $\ker \alpha_t$, we know that there is a unique solution Y^t by the contact condition. Notice that $Y^t_p = 0$. Thus we have found an X^t satisfying $\mathcal{L}_{X^t}\alpha_t + \dot{\alpha}_t = 0$.

We would now like to take the flow of X^t, but since this vector field is not complete this flow is not defined as an isotopy of U. However, given any $q \in U$ we can still define the curve $\varphi_{X^t}(q) \in U$ which exists for some time $t \in [0, T_q)$. Since $X^t_p = 0$ for all t the flow $\varphi_{X^t}(p)$ is constant and therefore it is defined for all $t \in [0,1]$. The set $\{(t,q); t < T_q\} \subseteq [0,1] \times U$ describing the times where the flow is defined is an open set, therefore there is some set $U' \subseteq U$ so that $\varphi_{X^t} : U' \to U$ is defined for all $t \in [0,1]$. We then let $\psi = f \circ \varphi_{X^1} : U' \to \mathbb{R}^{2n+1}$. □

Remark 4.15. In order to make the proof work for contact forms, we relied on the fact that α_t had no closed R_{α_t}-orbits, which we achieved by shrinking the set U (working locally). Thus while we can prove a Darboux theorem for contact forms, the proof cannot be extended to reprove Gray's theorem (which is a global theorem). Indeed, **Gray's theorem for contact forms is false**.

Note. For any constant C, the map $(x_i, y_i, z) \mapsto (Cx_i, Cy_i, C^2 z)$ is a contactomorphism of $\mathbb{R}^{2n+1}_{\text{std}}$. This says that $\mathbb{R}^{2n+1}_{\text{std}} = \bigcup U_i$, where each U_i is contactomorphic to the unit ball in $\mathbb{R}^{2n+1}_{\text{std}}$, and $U_i \subseteq U_{i+1}$. A "telescoping" argument then shows that $\mathbb{R}^{2n+1}_{\text{std}}$ is contactmorphic to the unit ball $B^{2n+1} \subseteq \mathbb{R}^{2n+1}_{\text{std}}$. Combining this observation with Darboux's theorem tells us that *around every point in a contact manifold, there is a neighborhood which is contactomorphic to $\mathbb{R}^{2n+1}_{std}$*.

5 Symplectic geometry

5.1 The basics

Symplectic geometry is an even dimensional geometry, intimately connected with contact geometry. Not only are these connections of abstract interest, they also give us a wealth of examples in contact geometry.

Definition 5.1. Let W^{2n} be a smooth manifold. A *symplectic structure* on W is a 2-form $\omega \in \Omega^2 W$ which is both closed and non-degenerate: $d\omega = 0$ and $\omega^{\wedge n}$ is never vanishing. A diffeomorphism between symplectic manifolds $\varphi : (W_1, \omega_1) \to (W_2, \omega_2)$ is called a *symplectomorphism* if $\varphi^*(\omega_2) = \omega_1$.

Example 5.2. Given any smooth manifold Q, we claim that the total space T^*Q is canonically symplectic. First we claim there is a 1-form $\lambda_{\text{std}} \in \Omega^1 T^*Q$ defined by the property that for any section $\sigma : Q \to T^*Q$, $\sigma^*\lambda_{\text{std}} = \sigma$ (make sure this formula makes sense to you). Together with the condition $\lambda_{\text{std}}|_{\{\text{fibers}\}} = 0$ guarantees the uniqueness of such a λ_{std}. To check existence we define λ_{std} in coordinates. Given any local coordinates $(q_1, \dots, q_n)$ on Q, let $(p_1, \dots, p_n)$ be dual coordinates on the fiber T^*Q_q, defined by $p_j(\partial_{q_j}) = p_j$ and $p_i(\partial_{q_j}) = 0$ for $i \neq j$. Then $\lambda_{\text{std}} = \sum\limits_{j=1}^{n} p_j dq_j$. Check that λ_{std} has the property described above, showing that it is independent of coordinates and therefore glues to a global 1-form.

We then define $\omega_{\text{std}} = d\lambda_{\text{std}}$. ω_{std} is closed because it is exact, and it is non-degenerate because in coordinates it is equal to $\sum\limits_{j=1}^{n} dp_j \wedge dq_j$. (Note that many authors instead use $-\omega_{\text{std}}$ as the canonical symplectic form.)

The previous example is a particular type of symplectic manifold where ω is an exact 2-form. We call such symplectic manifolds *exact*. Note that a compact manifold cannot have an exact symplectic structure. Indeed if W is compact, $\omega^{\wedge n} \in \Omega^{2n} W$ never being zero implies that $[\omega^{\wedge n}] \neq 0 \in H^{2n}_{dR} W \cong \mathbb{R}$, which implies that $[\omega] \neq 0 \in H^2 W$ since $\wedge$ descends to an operation on homology (cup product).

Definition 5.3. Let W be a complex manifold. A symplectic structure ω on W is called *Kähler* if the function $(v, w) \mapsto \omega(v, \sqrt{-1}w)$ is a Riemannian metric.

A nice property of Kähler manifolds is that complex submanifolds are automatically Kähler (this is immediate to prove). This gives us a wealth of examples of symplectic manifolds. Indeed, $(\mathbb{C}^n, \omega_{\text{std}} = \sum_j dx_j \wedge dy_j)$ is Kähler, so the zero set of any holomorphic functions on $\mathbb{C}^n$ is symplectic (as long as it is smooth). Similarly, $\mathbb{CP}^n$ is Kähler, with the Fubini-Study symplectic form. This implies that any smooth complex projective variety is symplectic.

Definition 5.4. Let (W, ω) be a symplectic manifold. A vector field X on W is called a *symplectic vector field* if its flow generates an isotopy through symplectomorphisms (a *symplectic isotopy*). Equivalently, X is a symplectic vector field if and only if $\mathcal{L}_X\omega = 0$.

Proposition 5.5. *Let (W, ω) be symplectic; non-degeneracy implies that $\omega^* : \Gamma(TW) \to \Omega^1 W$ is an isomorphism. Under this isomorphism, symplectic vector fields are identified with closed 1-forms.*

Proof: This is immediate from Cartan's formula: X is symplectic iff $0 = \mathcal{L}_X\omega = X \lrcorner\, d\omega + d(X \lrcorner\, \omega) = d(X \lrcorner\, \omega)$. □

In some ways the symplectic world is not as nice as the contact case: there are often cohomological conditions which does not appear in the contact world. For example, the isotopy extension theorem 4.6 does not hold in the symplectic world. The reason is simple: closed 1-forms do not admit cutoffs. For example, $d\theta$ is a closed 1-form on the annulus $A = \{a < x^2 + y^2 < b\} \subseteq \mathbb{R}^2$, but it cannot extend to a closed 1-form on $\mathbb{R}^2$. Indeed, any closed 1-form on $\mathbb{R}^2$ is exact, but $[d\theta] \neq 0 \in H^1(A)$. If we choose a symplectic form ω on $\mathbb{R}^2$ then the flow of the symplectic vector field $(\omega^*)^{-1}(d\theta)$ gives a path of symplectic embeddings $A \to \mathbb{R}^2$ which doesn't extend to a symplectic isotopy of $\mathbb{R}^2$.

To counter this, we define *Hamiltonian vector fields* to be symplectic vector fields X so that $X \lrcorner\, \omega$ is exact. Then Hamiltonian vector fields do admit cutoff functions.

For the purpose of relating symplectic geometry to the contact world, we are interested in a different type of vector field.

Definition 5.6. Let (W, ω) be a symplectic manifold. A *Liouville vector field* is a vector field Y so that $\mathcal{L}_Y\omega = \omega$.

Liouville vector fields are not symplectic, so the flow of a Liouville vector field doesn't preserve ω, but it does preserve it *conformally.* Indeed, if φ_t is the flow of a Liouville vector field, then $\varphi_t^*(\omega) = e^t\omega$.

ω^* identifies Liouville vector fields on W with primitives of ω: $\omega = \mathcal{L}_Y\omega = d(Y \lrcorner\, \omega)$. Therefore Liouville vector fields can only exist on exact symplectic manifolds, and a choice of λ satisfying $d\lambda = \omega$ gives a Liouville vector field $Y = (\omega^*)^{-1}(\lambda)$.

Notice that it is possible for Y to be zero on a subset of W, even though $\mathcal{L}_Y\omega \neq 0$ (since $\mathcal{L}_Y\omega$ is not tensorial in Y). But $Y \lrcorner\, \omega$ is tensorial in Y, therefore the set where Y vanishes is identical to the set where λ vanishes.

5.2 Contact and symplectic hypersurfaces

Proposition 5.7. *Let (W, ω) be a symplectic manifold, and let $M \subseteq W$ be a hypersurface (a codimension 1 submanifold). Suppose there is a Liouville vector field Y defined in a neighborhood of M, so that Y is transverse to M everywhere. Then $\ker(Y \lrcorner\, \omega|_M)$ is a contact structure on M.*

Proof: Let $\lambda = Y \lrcorner\, \omega$, so $d\lambda = \omega$. Along M, $TW = TM \oplus \langle Y \rangle$, and therefore $\omega^{\wedge n}$ being non-vanishing implies that $(Y \lrcorner\, \omega^{\wedge n})|_M$ is non-vanishing. But $Y \lrcorner\, \omega^{\wedge n} = n(Y \lrcorner\, \omega) \wedge \omega^{\wedge n-1} = n\lambda \wedge (d\lambda)^{\wedge n-1}$. □

This proposition gives us a wealth of compact contact manifolds:

Example 5.8. On $(T^*Q, \omega_{\text{std}})$, the primitive λ_{std} defines a Liouville vector field Y which is radial expansion along the fibers (in coordinates, $Y = \sum_j p_j \partial_{p_j}$). Choose a metric g on Q, and let $S^*Q = \{(q,p) \in T^*Q; g^*(p,p) = 1\}$ be the unit sphere bundle in the cotangent space. Then $\alpha_g = \lambda_{\text{std}}|_{S^*Q}$ is a contact form. Each cosphere fiber S^*Q_q is tangent to $\ker \alpha_g$. One can check that the flow of R_{α_g} is equal to the geodesic flow of the metric g, after identifying S^*Q with the unit sphere bundle in the tangent bundle, using g.

Notice that S^*Q is always cooriented, even if Q is non-orientable. On the other hand, $\mathbb{P}(T^*Q) = S^*Q/(q,p) \sim (q,-p)$ is never coorientable, even if Q is oriented. The contact structures $\ker \alpha_1$ and $\ker \alpha_{1/2}$ on T^3 given in Example 3.3 are contactomorphic to S^*T^2 and $\mathbb{P}(T^*T^2)$, respectively.

Example 5.9. Let W be any complex submanifold of $\mathbb{C}^N$, and let $M = S^{2N-1} \cap W$, where S^{2N-1} is a round sphere centered at the origin, with radius chosen so that the intersection with W is transverse. Then we claim that M is contact. Since W is complex $\omega = \omega_{\text{std}}|_W$ is a symplectic structure. Let $\lambda = \frac{1}{2}\sum_j x_j dy_j - y_j dx_j|_W$. Notice that $\lambda(\sqrt{-1}X) = \frac{1}{2}(\sum_j x_j dx_j + y_j dy_j)(X) = \frac{1}{4}X(\sum_j x_j^2 + y_j^2)$ for any vector X. λ is a primitive of ω, so it defines a Liouville vector field Y on W. Since S^{2N-1} is transverse to W, the form $\sum_j x_j dx_j + y_j dy_j|_W$ is never zero along M, which implies that λ is also non-vanishing, and thus Y is nonzero along M. The Kähler condition then implies that at each point, $0 < \omega(Y, \sqrt{-1}Y) = (Y \lrcorner \omega)(\sqrt{-1}Y) = \lambda(\sqrt{-1}Y) = \frac{1}{4}Y(\sum_j x_j^2 + y_j^2)$. Therefore $Y \pitchfork S^{2N-1}$, meaning $Y \pitchfork M$, so $\lambda|_M$ is a contact form.

We can describe the contact structure in another way. Suppose $X \in TM$, then as above $\lambda(X) = \frac{1}{2}(\sum_j x_j dx_j + y_j dy_j)(-\sqrt{-1}X)$. Therefore, $X \in \ker \lambda|_M = \xi$ if and only if $-\sqrt{-1}X$ is tangent to S^{2N-1}. $-\sqrt{-1}X$ is always tangent to W (since W is a complex submanifold), therefore $X \in \xi$ iff $-\sqrt{-1}X \in TM$. Therefore $\xi = TM \cap \sqrt{-1}TM$.

The previous proposition says that a hypersurface in a symplectic manifold is contact, given that it is transverse to a Liouville vector field. Conversely, we show that hypersurfaces in contact manifolds are exact symplectic manifolds, given a transversality condition.

Proposition 5.10. *Let (M, ξ) be a contact manifold, and let $W \subseteq M$ be a hypersurface. Suppose there is a contact vector field X defined near W, so that X is transverse to W, and never contained in ξ. Then W inherits a canonical exact symplectic structure (with a canonical Liouville vector field).*

Proof: Given such an X, we can define a contact form α for ξ by defining $\alpha|_\xi = 0$, $\alpha(X) = 1$ (thus $R_\alpha = X$). Since $TM|_W = TW \oplus \langle X \rangle$, the fact that $\alpha \wedge (d\alpha)^{\wedge n} \neq 0$ implies that $(d\alpha)^{\wedge n}|_W = \alpha(X)(d\alpha)^{\wedge n} = X \lrcorner \alpha \wedge (d\alpha)^{\wedge n} \neq 0$. Thus $\omega = d\alpha|_W$ is symplectic, and $\lambda = \alpha|_W$ is a primitive of ω (which in turn defines a Liouville vector field). □

Generally, this proposition is less useful than the previous one, for two reasons. First of all, since exact symplectic structures only exist on open manifolds, we know that the previous proposition only applies when W is an open hypersurface, which is a major drawback. Secondly, in terms of generating examples, the earlier proposition is far more useful since it allows us to construct contact manifolds from symplectic manifolds, which in turn can be constructed from holomorphic manifolds. But since we are generally short on contact examples, constructing symplectic manifolds from contact manifolds tends to not give us anything new.

5.3 Symplectizations and contactizations

We now investigate another way in which contact and symplectic geometry intertwine: rather than looking at hypersurfaces, we look at $\mathbb{R}$-equivariant structures.

Proposition/Definition 5.11. Let ξ be a hyperplane distribution on a $(2n-1)$-dimensional manifold M. Let $S(M,\xi) = \{(q,p); q \in M, p \in T^*M_q, \ker p = \xi_q\} \subseteq T^*M$. Then $\omega_{\text{std}}|_{S(M,\xi)}$ is a symplectic structure if and only if ξ is contact. In this case, we call the symplectic manifold $S(M,\xi)$ the *full symplectization* of (M,ξ).

Proof: The statement is local in M, so we can choose $\alpha \in \Omega^1 M$ with $\ker\alpha = \xi$. Then for any $(q,p) \in S(M,\xi)$, $p = r\alpha_q$ for some $r \in \mathbb{R}^* = \mathbb{R}\setminus\{0\}$. Thus $S(M,\xi)$ is diffeomorphic to $M \times \mathbb{R}^*$. (This is only true locally. Globally, $S(M,\xi)$ is an $\mathbb{R}^*$ bundle over M, bundle isomorphic to $(TM/\xi) \setminus \{\text{zero section}\}$.)

We claim that $\lambda_{\text{std}}|_{S(M,\xi)} = r\alpha$ with this coordinate identification. The easiest way to check this is simply to notice that is $\sigma : M \to S(M,\xi)$ is a section, then $\sigma = f\alpha$ for some non-vanishing function $f \in C^\infty M$, and $\sigma = f\alpha = \sigma^*(r\alpha)$. Since this is the defining property of λ_{std} it follows that they are equal. But then $\omega_{\text{std}}|_{S(M,\xi)} = d(r\alpha) = dr\wedge\alpha + rd\alpha$, and so $(\omega_{\text{std}}|_{S(M,\xi)})^{\wedge n} = r^{n-1}dr\wedge\alpha\wedge(d\alpha)^{\wedge n-1}$. The result follows. □

If (M,ξ) is not coorientable, then $S(M,\xi)$ is connected. However if ξ has a coorientation, then $S(M,\xi)$ is the disjoint union of two components $S^+(M,\xi)$ and $S^-(M,\xi)$. If (M,ξ) is a cooriented contact manifold, we define the *positive symplectization* to be the symplectic manifold $S^+(M,\xi)$. Notice that $S^+(M,\xi)$ and $S^-(M,\xi)$ are symplectomorphic, so if ξ is only coorientable (but without a given coorientation) $S^+(M,\xi)$ is still a well defined symplectic manifold, up to symplectomorphism. If we simply say the *symplectization* of a contact manifold (M,ξ), we will assume that we are talking about the positive symplectization whenever ξ is coorientable, and the full symplectization if not.

Supposing that (M,ξ) is coorientable, a section $\alpha : M \to S(M,\xi)$ is the same as choosing a contact form for ξ, by definition of $S(M,\xi)$. This also be seen in the following way. A section α embeds M into $S(M,\xi)$ as a hypersurface which is transverse to the Liouville vector field. Then Proposition 5.10, $\lambda_{\text{std}}|_{\alpha(M)}$ is a contact form on M. But by the canonical property of λ_{std}, $\lambda_{\text{std}}|_{\alpha(M)} = \alpha^*\lambda_{\text{std}} = \alpha$.

In the presence of a contact form, we can write $S^+(M,\ker\alpha,\lambda_{\text{std}}) \cong (\mathbb{R}_+ \times M, r\alpha)$. This is because $S^+(M,\ker\alpha)$ is a $\mathbb{R}_+$ bundle over M, so the section $\alpha : M \to S^+(M, \ker\alpha)$ defines a diffeomorphism $S^+(M, \ker\alpha) \cong \mathbb{R}_+ \times M$, so that $\alpha(M) = \{1\} \times M$. Notice $\mathcal{L}_Y \lambda_{\text{std}} = Y \lrcorner d\lambda_{\text{std}} + d(\lambda_{\text{std}}(Y)) = Y \lrcorner \omega_{\text{std}} + d(\omega_{\text{std}}(Y, Y)) = \lambda_{\text{std}}$. So $\mathcal{L}_{r\partial_r}\lambda_{\text{std}} = \lambda_{\text{std}}$ and $\lambda_{\text{std}}|_{\{1\}\times M} = \alpha$; interpreting this as a differential equation with given initial conditions it follows that $\lambda_{\text{std}} = r\alpha$.

Because the symplectization is defined as a submanifold of T^*M, it comes not only with a symplectic structure, but also a Liouville structure $\lambda_{\text{std}}|_{S(M,\xi)}$, which defines a Liouville vector field Y (in the coordinates above, $Y = r\partial_r$). Y is a complete vector field (in both directions), therefore its flow induces an action of $\mathbb{R}$ on $S(M, \xi)$. Conversely,

Proposition 5.12. *Let (W, ω) be a connected symplectic manifold with a Liouville $\mathbb{R}$-action, that is, for any $t \in \mathbb{R}$ we get a diffeomorphism $\varphi_t : W \to W$ satisfying $\varphi_t^*\omega = e^t\omega$. Suppose the quotient $M = W/\mathbb{R}$ is a smooth manifold. Then (W, ω) is a positive symplectization of a contact manifold (M, ξ), uniquely defined up to contactomorphism.*

Proof: Let $Y = \frac{d\varphi_t}{dt}|_{t=0}$, by assumption Y is a Liouville vector field on W, and therefore $\lambda = Y \lrcorner \omega$ satisfies $d\lambda = \omega$. Then $\mathcal{L}_Y\lambda = Y \lrcorner d\lambda + d(\lambda(Y)) = Y \lrcorner \omega + d(\omega(Y, Y)) = \lambda + 0$. Since $\mathbb{R}$ is contractible it follows that W is diffeomorphic to $M \times \mathbb{R}$, by a diffeomorphism sending Y to ∂_t, where t is the coordinate in $\mathbb{R}$. This diffeomorphism is not canonical, but we choose one arbitrarily.

Let $\alpha = \lambda|_{M\times\{0\}}$. $\mathcal{L}_{\partial_t}\lambda = \lambda$ implies that $\phi_t^*\lambda = e^t\lambda$, which implies that $\lambda = e^t(\alpha + f dt)$ for some function $f \in C^\infty M$. But since $\lambda(\partial_t) = \lambda(Y) = 0$, the function f is identically zero. Therefore $\lambda = e^t\alpha$, so the coordinate change $r = e^t$ identifies W with the symplectization. □

Example 5.13. In the previous section, we defined the contact manifold $S^*Q \subseteq T^*Q$ by choosing a metric on Q. We give an alternative viewpoint of this contact manifold by defining it in a metric invariant way. Notice that on T^*Q, the Liouville form λ_{std} is zero exactly along the zero section $Z = \{p = 0\} \subseteq T^*Q$. The Liouville flow $\varphi_t : T^*Q \to T^*Q$ acts by $\varphi_t(q, p) = (q, e^t p)$. Therefore $p_t|_{T^*Q\setminus Z}$ defines a free action, and its quotient is diffeomorphic to a manifold. Indeed it is diffeomorphic to S^*Q, since a choice of metric of Q defines an embedding $S^*Q \hookrightarrow T^*Q \setminus Z$ which intersects each Liouville orbit exactly once. So $(T^*Q \setminus Z, \lambda_{\text{std}})$ is a positive symplectization of some contact structure on S^*Q, but since the natural contact structure $\ker\alpha_g$ on S^*Q is *defined* as $\alpha_g = \lambda_{\text{std}}|_{S^*Q}$, we see that $(T^*Q \setminus Z, \lambda_{\text{std}}) = C(S^*Q, \ker\alpha_g)$.

For yet another viewpoint, we define $\mathbb{P}^+(T^*Q) = \{(q, P); P \text{ is a cooriented hyperplane at the point } q \in Q\}$. We define a contact structure ξ_{std} on $\mathbb{P}^+T^*Q$ as follows. At a point $(q, P) \in \mathbb{P}^+T^*Q$, a vector (∂_q, ∂_P) is contained in ξ_{std} iff $\partial_q \in P$. Perhaps the easiest way to show that ξ_{std} is contact is by showing that $\ker : S^*Q \to \mathbb{P}^+T^*Q$ is a contactomorphism from $\ker\alpha_g$ to ξ_{std} (for any choice of metric g). Given a point $(q, p) \in S^*Q$, $(\alpha_g)_{(q,p)}(\partial_q, \partial_p) = p(\partial_q)$ (since α_g is the restriction of $\lambda_{\text{std}} \in \Omega^1T^*Q$). Therefore $(\partial_q, \partial_p) \in \ker\alpha_g$ iff $\partial_q \in \ker p = P$.

We also note that the identification in Proposition 5.12 is functorial: contact isotopies correspond to $\mathbb{R}$-equivariant Hamiltonian isotopies. Let (M, ξ) be a coorientable contact manifold, and let Y be the canonical Liouville vector field on $W = S^+(M,\xi)$. Then sections $h : M \to TM/\xi$ are in bijective correspondence with functions $H : W \to \mathbb{R}$ satisfying $Y(H) = H$. Indeed, any function H of this form can be written as $H(q,p) = p(h(q))$, giving the correspondence.

The function H defines a Hamiltonian vector field $X_H = (\omega^*)^{-1}(dH)$ on W, and the function h defines a contact vector field X_h on M. Since $dH = X_H \lrcorner\, \omega$, we see

$$\mathcal{L}_Y dH = \mathcal{L}_Y(X_H \lrcorner\, \omega)$$
$$Y \lrcorner\, d^2H + d(dH(Y)) = [Y, X_H] \lrcorner\, \omega + X_H \lrcorner\, (\mathcal{L}_Y \omega)$$
$$d(H) = [Y, X_H] \lrcorner\, \omega + X_H \lrcorner\, \omega$$

which implies that $[Y, X_H] \lrcorner\, \omega = 0$, so $[Y, X_H] = 0$ by non-degeneracy of ω. Therefore X_H projects to a vector field on M, we claim that it projects to X_h. (Since this is a local statement, it can be checked by calculating with a contact form on M, or even by choosing Darboux coordinates.)

Note. This correspondence of contact and Hamiltonian flows also holds for the full symplectization of a contact manifold, with a slight restatement: the condition $Y(H) = H$ is not enough to guarantee that $H(q,p) = p(h(q))$. We must further assume that $H(q,-p) = -H(q,p)$, with this additional condition everything above holds for the full symplectization.

We now define a similar construction, in the opposite direction.

Proposition/Definition 5.14. Let W be an even dimensional manifold, and let $\lambda \in \Omega^1(W)$. Let $C(W,\lambda)$ be the manifold $W \times \mathbb{R}$, equipped with the 1-form $\alpha_{\text{std}} = dz - \lambda$ ($z \in \mathbb{R}$). Then $d\lambda$ is symplectic if and only if $\ker \alpha_{\text{std}}$ is contact. In this case, we call $C(W,\lambda)$ the *contactization* of (W,λ).

In a symplectization, any section $M \to S(M,\xi)$ defines a contact form by restricting λ_{std}. Similarly, the restriction of α_{std} to a section $W \to C(W,\lambda)$ defines a 1-form $\widetilde{\lambda} = \lambda + df$ satisfying $d\widetilde{\lambda} = d\lambda$ (it chooses a primitive for the symplectic structure $d\lambda$). Notice that $R_{\alpha_{\text{std}}} = \partial_z$ for any contactization.

Proposition 5.15. *Let $(M, \ker\alpha)$ be a contact manifold, then the flow φ_t of R_α defines an action on M of the group $\mathbb{R}$. Assume that this action has no fixed points, and that the quotient $M/\mathbb{R}$ is a smooth manifold. Then (M,ξ) is contactomorphic to $C(W,\lambda)$ for some exact symplectic manifold (W,λ).*

Similar to the symplectization case, this identification is also functorial with respect to flows. If $f \in C^\infty W \times \mathbb{R}$ is a function satisfying $R_{\alpha_{\text{std}}}(f) = 0$, then f descends to a function of W, and the contact flow of f on $C(W,\lambda)$ projects to the Hamiltonian flow of f on W.

An important example of a contactization is the contactization of a cotagent bundle, which we call the *first jet bundle of* Q, $(\mathcal{J}^1Q, \alpha_{\text{std}}) = C(T^*Q, \lambda_{\text{std}}) = (T^*Q \times \mathbb{R}, dz - \sum p_i dq_i)$. Notice that $\mathbb{R}^{2n+1}_{\text{std}} = \mathcal{J}^1\mathbb{R}^n$. We will see the significance of this example in the next section.

6 Legendrian submanifolds

Definition 6.1. Let (M, ξ) be a contact manifold. A submanifold $\Lambda \subseteq M$ is called *isotropic* if it is tangent to ξ, that is $T\Lambda_q \subseteq \xi_q$ for all $q \in \Lambda$.

Of course, since ξ is a non-integrable plane field it is not possible for $T\Lambda_q = \xi_q$; the dimension of any isotropic submanifold must be less than the dimension of ξ. Since Λ is a submanifold, we know that any two vector fields $X, Y \in \Gamma(T\Lambda)$ which are tangent to Λ have a Lie bracket which is also tangent to Λ: $[X, Y] \in \Gamma(T\Lambda)$. Therefore, $\omega_\xi|\Lambda = 0$. Proposition 2.4 immediately tells us

Proposition/Definition 6.2. Let (M^{2n+1}, ξ) be a contact manifold, and let $\Lambda \subseteq M$ be an isotropic submanifold. Then $\dim \Lambda \leqslant n$. If $\dim \Lambda = n$ we say that Λ is *Legendrian.*

If $\xi = \ker \alpha$, then of course $T\Lambda \subseteq \xi$ iff $\alpha|_\Lambda = 0$. This of course implies that $d\alpha|_\Lambda = 0$, which is the same as the above claim $\omega_\xi|_\Lambda = 0$.

Example 6.3. Let Q be any smooth manifold, and let $\Lambda \subseteq \mathcal{J}^1 Q$ be an n-dimensional submanifold so that the projection $\pi : \mathcal{J}^1 Q \to Q$ is a diffeomorphism when restricted to Λ. This condition is equivalent to Λ being the image of a section $\sigma : Q \to \mathcal{J}^1 Q$. Since $\mathcal{J}^1 Q \simeq T^*Q \times \mathbb{R}$, we can write $\sigma = (\beta, f)$, where $\beta \in \Omega^1 Q$ and $f \in C^\infty Q$. Then $\alpha_{\mathrm{std}}|_\Lambda = \pi^*\sigma^*\alpha_{\mathrm{std}} = \pi^*(\beta, f)^*(dz - \lambda_{\mathrm{std}}) = \pi^*(f^*(dz) - \beta^*(\lambda_{\mathrm{std}})) = \pi^*(df - \beta)$. Of course π^* is an isomorphism, therefore we see that $\Lambda = \sigma(Q)$ *is Legendrian iff* $df = \beta$.

A nice interpretation of this perspective is reinterpreting first-order differential equations on Q. Such a differential equation can be written as $G(q, f(q), df_q) = 0$, and we are trying to understand the space of functions $f \in C^\infty Q$ which satisfy the equation. Then, G should be interpreted simply as a function $G \in C^\infty \mathcal{J}^1 Q$. And solutions to the differential equation G are sections $\sigma : Q \to \mathcal{J}^1 Q$ which satisfy two properties: $G \circ \sigma = 0$, and $\sigma(Q)$ is Legendrian. Notice that the first condition is totally flexible: finding sections σ with $\sigma(Q) \subseteq G^{-1}(0)$ is a question purely in the realm of differential topology. In this sense, all first order differential equations can be reduced to a universal differential equation $\sigma^*\alpha_{\mathrm{std}} = 0$.

Note. Though this business of reducing all differential equations to a Legendrian condition may seem powerful, in fact it is something that any first-semester student of ODEs knows how to do: given a differential equation $G(x, y, \dot{y}) = 0$, we can reduce it to the algebraic equation $G(x, y, u) = 0$ and the "universal differential equation" $u = \dot{y}$. This language of 1-jet spaces allows us to do the same thing for first order PDEs on manifolds Q and to rephrase the differential condition more geometrically (being tangent to ξ_{std}), but the general idea is exactly the same.

There are higher jet spaces $\mathcal{J}^k Q$ which have identical properties for higher order differential equations, but they do not have any natural contact structure. The natural structure on $\mathcal{J}^k Q$ is a non-integrable distribution of higher codimension, but these plane fields do not satisfy the same nice geometric properties as contact structures (Gray stability, Darboux theorems).

Example 6.4. Let $K \subseteq Q$ be any submanifold with positive codimension. We define a Legendrian $\nu^*K \subseteq S^*Q$, called the *conormal bundle* of K. $\nu^*K = \{(q,p) \in S^*Q; q \in K \text{ and } p(TK_q) = 0\}$. (In $\mathbb{P}^+T^*Q$, this corresponds to $\nu^*K = \{(q,P) \in \mathbb{P}^+T^*Q; q \in K \text{ and } TK_q \subseteq P\}$). Check that ν^*K is Legendrian. Notice that ν^*K is well defined for any *immersed* submanifold K, and as long as the immersion is transverse to itself then ν^*K will be an embedded Legendrian.

We point out two extreme cases of this. When K is a point in Q, ν^*K is just the fiber of S^*Q. If K is a hypersurface, then ν^*K is a two-to-one cover of K, since $\nu^*K = \{(q,p); \ker p = TK_q\}$. It is the coorientation double cover of K. If K is cooriented in Q then this cover is disconnected, and we can choose a canonical component $\nu^*_+K = \{(q,p); \ker p = TK_q \text{ as cooriented hyperplanes}\}$.

Note. By choosing a metric on Q we get a contact form α_g on S^*Q, which gives rise to the Reeb flow $\varphi_t : S^*Q \to S^*Q$. We noted in Example 5.8 that φ_t is equal to the geodesic flow, using the identification under the diffeomorphism $g^* : SQ \to S^*Q$ (where $SQ = \{(q,v) \in TQ; g(v,v) = 1\}$).

Given a cooriented hypersurface $K \subseteq Q$, φ_t acts on ν^*_+K by flowing each point of K in the direction orthogonal to K, at unit speed. (This is not the same as mean curvature flow, which also flows hypersurfaces orthogonally; the speed of the flow at a point is determined by the mean curvature. For instance, the mean curvature flow acting on a round circle in flat $\mathbb{R}^2$ will shrink the radius at a rate $r = r_0 - t^2$, whereas φ_t will shrink the radius of an inwardly cooriented circle linearly in time.)

Because φ_t has this unit speed property, it is very useful is geometric optics for modeling the time evolution of wave fronts of light. In fact, this is one of the earliest motivations for developing contact geometry. Suppose the point $q \in Q$ is to be thought of as a point source which is radiating light. We think of our initial wavefront as the fiber $\Lambda_0 = S^*Q_q$: all the light in concentrated at q and it is moving in every direction. Then $\varphi_t(\Lambda_0)$ is the time evolution of the wavefront of the light. For small t we have $\varphi_t(\Lambda_0) = \nu^*_+S^{n-1}(q,t)$, where $S^{n-1}(q,t)$ is the sphere of radius t centered at q (outwardly cooriented). But for larger t the wavefront $\pi(\varphi_t(\Lambda_0))$ will have interesting geometry, self-intersections, and singularities. (Here $\pi : S^*Q \to Q$ is the projection. Of course $\varphi_t(\Lambda_0)$ will be a smooth embedded Legendrian for all $t \in \mathbb{R}$, since φ_t is a contact isotopy.)

The next theorem should be thought of as a strengthening of the Darboux theorem.

Theorem 6.5 (Legendrian Neighborhood Theorem)**.** *Let $\Lambda \subseteq (M,\xi)$ be an embedded Legendrian submanifold, and assume that $\xi|_\Lambda \subseteq TM|_\Lambda$ is coorientable. Then there is an open neighborhood U of Λ and a contact embedding $\psi : U \to \mathcal{J}^1\Lambda$, so that $\psi(\Lambda)$ is the zero section $Z = \{z = 0, p = 0\} \subseteq \mathcal{J}^1\Lambda$. Furthermore, given $\alpha \in \Omega^1 M$ satisfying $\ker \alpha = \xi$, we can choose U and ψ so that $\psi^*\alpha_{std} = \alpha$.*

Proof: The latter statement ($\psi^*\alpha_{\text{std}} = \alpha$) implies the former ($\psi_*\xi = \ker \alpha_{\text{std}}$). First, we show that there is a bundle isomorphism $f : TM|_\Lambda \to T\mathcal{J}^1\Lambda|_Z$, so that $f(T\Lambda) = TZ$ and $f(\xi|_\Lambda) = \ker \alpha_{\text{std}}|_\Lambda$. The diffeomorphism $\mathcal{J}^1\Lambda \simeq T^*\Lambda \times \mathbb{R}$ gives a canonical bundle isomorphism $T\mathcal{J}^1\Lambda|_Z \cong T\Lambda \oplus T^*\Lambda \oplus \mathbb{R}$, identifying

TZ with $T\Lambda$ and $\ker \alpha_{\mathrm{std}}|_\Lambda$ with $T\Lambda \oplus T^*\Lambda$. By assumption we can choose $\alpha \in \Omega^1\nu$ with $\ker\alpha = \xi|_\nu$. We let $f|_{T\Lambda}$ be the identity map to $TZ = T\Lambda$, and define $f(R_\alpha) = \partial_z$. A vector $v \in \xi_q$ defines a covector $(v \lrcorner d\alpha_q)|_{T\Lambda_q}$, and furthermore this covector is zero if and only if $v \in T\Lambda$. Therefore, $\cdot \lrcorner d\alpha$ defines an isomorphism from $\xi|_\Lambda/T\Lambda$ to $T^*\Lambda$, this completes the construction of f.

Being a normal bundle, we can identify U diffeomorphically with $TM|_\Lambda/T\Lambda$. Composing with f, we get a diffeomorphism $g : U \to T\mathcal{J}^1\Lambda|_Z/TZ \cong \mathcal{J}^1\Lambda$. This diffeomorphism has the property $g(\Lambda) = Z$, and also $g^*(\alpha_{\mathrm{std},q}) = \alpha_q$ for all $q \in \Lambda$.

We now run a standard Gray stability argument. We have two contact forms on U, α and $g^*\alpha_{\mathrm{std}}$. For $t \in [0,1]$, let $\alpha_t = tg^*\alpha_{\mathrm{std}} + (1-t)\alpha$. When restricted to Λ α_t is constant, so α_t is contact along Λ. By shrinking U, we can therefore assume that $\ker\alpha_t$ is contact everywhere in U for all $t \in [0,1]$. By shrinking U further so that R_{α_t} has no closed orbits in U, we can find functions $h^t \in C^\infty U$ satisfying $R_{\alpha_t}(h^t) = -\dot\alpha_t(R_{\alpha_t})$, so that $h^t(\Lambda) = 0$ and $dh^t|_\Lambda = 0$ (compare with the proof of Theorem 4.14). We also find $Y^t \in \ker\alpha_t$ so that $Y^t \lrcorner d\alpha = -dh^t - \dot\alpha_t$.

Then $X^t = Y^t + h^t R_{\alpha_t}$ satisfies $\mathcal{L}_{X^t}\alpha_t = -\dot\alpha_t$ everywhere. $X^t = 0$ along Λ, so the flow of X^t is defined for all $t \in [0,1]$ along Λ. By choosing an open $U' \subseteq U$ we can define $\varphi_t : U' \to U$ to be the flow of X^t for all $t \in [0,1]$. Then $\varphi_t^*(\alpha_t) = \alpha$. Setting $\psi = g \circ \varphi_1$ completes the result. □

Note. The assumption that ξ is coorientable along Λ is essential: if ξ is defined on $S^1 \times D^2 = \{(\theta, x, y); \theta \in \mathbb{R}/\mathbb{Z}, (x,y) \in D^2\}$ by $\xi = \ker(\cos(\frac{\theta}{\pi})dy - \sin(\frac{\theta}{\pi})dx)$ then it's easy to check that no neighborhood of the Legendrian $\Lambda = \{x = 0, y = 0\}$ is contactomorphic to $\mathcal{J}^1S^1$, since $(TM/\xi)|_\Lambda$ is not orientable. However, one can show that, given the line bundle isomorphism class of $(TM/\xi)|_\Lambda$ (which is classified by $H^1(\Lambda;\mathbb{Z}_2)$), a small tubular neighbrhood of Λ is uniquely determined.

Similarly, suppose we are given an isotropic manifold $K^k \subseteq (M^{2n+1}, \xi)$. Assuming that $\xi|_K$ is coorientable, a neighborhood of K is contactomorphic to a symplectic vector bundle $E^{2n-2k} \to J^1K$. The symplectic bundle isomorphism type of E is a local invariant of K, but note that this is an algebro-topological invariant (it is classified by homotopy classes of maps $K \to BU(2n-2k)$). As a general philosophy, *the bundle isomorphism type of the normal bundle to an isotropic submanifold determines the contactomorphism type of neighborhoods of that manifold.* The Darboux theorem is a special case of this, since a point is an isotropic submanifold and any vector bundle over a point is trivial.

7 Hypersurfaces in 3-dimensional contact manifolds, part I

One of the most powerful approaches to low dimensional topology (and geometric topology in general) is cut-and-paste techniques: given a manifold, we can hope to chop it up into pieces we understand, and then glue those pieces back together. In order to apply this in the contact world, we need a way to say which gluings are possible.

Before embarking into the world of hypersurfaces, we make a note about contact 3-manifolds and orientations. Any contact 3-manifold (M, ξ) is canonically oriented, by declaring the basis $(X, Y, -[X, Y])$ to be a positive orientation for any linearly independent vector fields $X, Y \in \Gamma(\xi)$ (why is this well defined?). Note that this works even when ξ is not coorientable. In the language of contact forms, notice that the sign of $\alpha \wedge d\alpha$ does not depend on the sign of α, and therefore even local contact forms define a global orientation.

Note. Though we will not really rely on this fact, we note that *any orientable 3-manifold is parallelizable*, that is, $TM \cong M \times \mathbb{R}^3$ (but this identification is not canonical). This is a modestly deep theorem, which we won't have a chance to prove. The difficult part is showing that any orientable 3-manifold is spin. The standard proof follows from the Wu formulas relating Steenrod squares with Stiefel-Whitney classes. Another proof due to Kirby works more geometrically, relying on the fact that all homology classes in $H_2(M; \mathbb{Z}_2)$ can be represented by embedded surfaces.

7.1 Characteristic foliations as germs of contact structure

Definition 7.1. Let Σ^2 be a surface in a three manifold (M, ξ) equipped with a distribution. The *characteristic foliation* of ξ, denoted $\mathcal{F}_\xi$, is the *singular* foliation on Σ defined by $T\Sigma \cap \xi$.

Notice that in higher dimensions, there is no reason for the distribution $T\Sigma \cap \xi$ to be integrable, even when it is non-singular. But of course any 1-dimensional distribution is integrable.

While we haven't defined the term singular foliation, we can take its meaning to be exactly what is needed for our definition to work: a singular foliation is a smooth hyperplane distribution $\xi \subseteq T\Sigma \oplus L$, with the equivalence relation $\xi_1 \sim \xi_2$ iff $A(\xi_1) = \xi_2$ for some bundle automorphism $A : T\Sigma \oplus L \to T\Sigma \oplus L$ satisfying $A|_{T\Sigma} = \mathrm{id}$ (here L is an arbitrary line bundle). Because this definition may feel abstract, we give an alternative characterization of the concept.

Proposition 7.2. *Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be coorientable singular foliations on an oriented surface Σ. Then $\mathcal{F}_j$ is an equivalence class of distribution 2-dimensional distribution $\xi_j \subseteq T\Sigma \oplus \mathbb{R}$. We can therefore write $\xi_j = \ker(\beta_j + f_j dt)$, where $\beta_j \in \Omega^1\Sigma$ and $f_j \in C^\infty\Sigma$. Notice that $\ker \beta_j = \mathcal{F}_j$ everywhere. Then $\mathcal{F}_1 = \mathcal{F}_2$ if and only if $\beta_1 = g\beta_2$ for some nonvanishing function $g \in C^\infty\Sigma$.*

Proof: Let $\mathcal{F}_j$ be the equivalence class of $\ker(\beta_j + f_j dt)$. A bundle automorphism $A : T\Sigma \oplus \mathbb{R} \to T\Sigma \oplus \mathbb{R}$ satisfying $A|_{T\Sigma} = \mathrm{id}$ can be written as $A(v, t) = (v + tX, ht)$ for some vector field $X \in \Gamma(T\Sigma)$ and nonvanishing $h \in C^\infty\Sigma$ (here $v \in T\Sigma, t \in \mathbb{R}$). So $A^*(\beta_2 + f_2 dt) = \beta_2 + (\beta_2(X) + hf_2)dt$. Then $\mathcal{F}_1 = \mathcal{F}_2$ iff there exists A so that $\ker(\beta_1 + f_1 dt) = \ker(\beta_2 + (\beta_2(X) + hf_2)dt)$ which happens if and only if $\beta_1 + f_1 dt = g\beta_2 + g(\beta_2(X) + hf_2)dt$. So $\mathcal{F}_1 = \mathcal{F}_2$ implies $\beta_1 = g\beta_2$.

Conversely, assuming $\beta_1 = g\beta_2$, we need to show that the equation $f_1 = g(\beta_2(X) + hf_2)$ admits a solution for X and h. First, let $C \subseteq \Sigma$ be the vanishing set of β_1. Then β_2 also vanishes, so f_1 and f_2 must both be non-vanishing (since

$\beta_j + f_j dt \neq 0$). Let U be an open neighborhood in Σ containing C, so that f_2 is non-vanishing on $\overline{U}$, and let $h|_U = \frac{f_1}{g f_2}$. Then extend h to Σ arbitrarily. It remains to find X so that $f_1 - ghf_2 = g\beta_2(X) = \beta_1(X)$. The function on the left hand side vanishes on $\overline{U}$, and β_1 is non-vanishing on $\Sigma \setminus C$, so we can clearly find such an X (supported in $\Sigma \setminus U$). □

This proposition is very useful in practice. Writing $\xi = \ker \alpha$ allows us to write $\beta = \alpha|_\Sigma$. Therefore, the singular foliation $\mathcal{F}_\xi$ is determined by the equivalence class of $\alpha|_\Sigma$ under scaling by non-zero functions.

Remark 7.3. It is generally **not** enough to assume that singular sets of $\mathcal{F}_1$ and $\mathcal{F}_2$ are equal and the non-singular sets are equivalent as foliations. For example, let $\beta_1 = xdx + ydy$ and $\beta_2 = (x^2 + y^2)\beta_1$ be defining 1-forms for singular foliations on D^2. Then $\ker \beta_1 = \ker \beta_2$ everywhere, but the singular foliations are not equivalent.

Theorem 7.4. *Let $\Sigma_j \subseteq (M_j, \xi_j)$ be embedded orientable surfaces inside contact 3-manifolds, $j = 1, 2$. Suppose there is a diffeomorphism $\psi : \Sigma_1 \to \Sigma_2$ satisfying $\psi_*(\mathcal{F}_{\xi_1}) = \mathcal{F}_{\xi_2}$. Then there are neighborhoods $U_j \supseteq \Sigma_j$ and a contactomorphism $\varphi : U_1 \to U_2$, so that $\varphi|_{\Sigma_1} = \psi$.*

Proof: $\psi_* \times \mathrm{id}$ is a bundle isomorphism $T\Sigma_1 \oplus \mathbb{R} \to T\Sigma_2 \oplus \mathbb{R}$. By choosing identifications $TM_j|_{\Sigma_j} \cong T\Sigma_j \oplus \mathbb{R}$ we get distributions $\xi_j|_{\Sigma_j} \subseteq T\Sigma_j \oplus \mathbb{R}$, which define the characteristic foliations $\mathcal{F}_{\xi_j}$. By assumption $\psi_*(\mathcal{F}_{\xi_1}) = \mathcal{F}_{\xi_2}$, which is to say there exists a bundle automorphism $A : T\Sigma_2 \oplus \mathbb{R} \to T\Sigma_2 \oplus \mathbb{R}$ so that $(A \circ (\psi_* \times \mathrm{id}))(\xi_1) = \xi_2$. Thus we get a bundle isomorphism $\widetilde{\psi} : TM_1|_{\Sigma_1} \to TM_2|_{\Sigma_2}$ covering the diffeomorphism $\psi : \Sigma_1 \to \Sigma_2$, so that $\widetilde{\psi}(\xi_1|_{\Sigma_1}) = \xi_2|_{\Sigma_2}$.

Neighborhoods $U_j \supseteq \Sigma_j$ are diffeomorphic to $TM_j|_{\Sigma_j}/T\Sigma_j$, therefore $\widetilde{\psi}$ defines a diffeomorphism $\widehat{\psi} : U_1 \to U_2$, satisfying $\widehat{\psi}|_{\Sigma_1} = \psi$ and $\widehat{\psi}_*((\xi_1)_q) = (\xi_2)_{\psi(q)}$ for all $q \in \Sigma_1$. From here we run a Gray stability argument. Let ξ_t be a homotopy of distributions on U_2 connecting ξ_2 and $\widehat{\psi}_*(\xi_1)$, which is fixed on Σ_2. Since it is fixed on Σ_2 and being contact is an open condition, we can assume that ξ_t is contact everywhere, by shrinking U_2. By the contact condition, we can find vector fields $X^t \in \Gamma(\xi_t)$ so that $\mathcal{L}_{X^t}\xi_t = -\dot{\xi}_t$. We can then find $U_2' \subseteq U_2$ so that the flow of X^t, $\varphi_t : U_2' \to U_2$ is defined for all t (since $X^t|_{\Sigma_2} = 0$), and $\varphi_{t*}(\xi_t) = \xi_2$. Then $\varphi = \varphi_1 \circ \widehat{\psi} : \widehat{\psi}^{-1}(U_2') \subseteq U_1 \to U_2$ is the desired contactomorphism. □

Remark 7.5. In the proof, we never once took a second derivative of ψ. Therefore if ψ is assumed only to be a C^1-diffeomorphism, the entire proof runs through, and we get a C^1 contactomorphism $\varphi : U_1 \to U_2$. Suppose we use this contactomorphism to glue C^∞ contact manifolds together, $M = (M_1, \xi_1) \cup_\psi (M_2, \xi_2)$. Then M is a C^1-smooth manifold with a C^1-smooth contact structure ξ.

But then, any C^1 manifold is C^1 diffeomorphic to a *unique* C^∞ manifold $\widetilde{M}$, and this diffeomorphism pushes forward ξ to a C^1 contact structure on $\widetilde{M}$. But then, if ζ_1 and ζ_2 are two C^∞ contact structures which are C^1 close to ξ, then they are C^1 close to each other, and therefore Corollary 4.10 implies that ζ_1 is contactomorphic to ζ_2.

The takeaway from all of this is: C^1 gluings of C^∞ contact manifolds are good enough to define C^∞ contact structures on C^∞ manifolds, canonical up to C^∞ contactomorphism. Though we do not study this in depth, this is actually very important for us, because it the C^1 classification of singularities of $\mathcal{F}_\xi$ is known, whereas the C^∞ classification is far more difficult.

Note. Notice that we assumed nowhere that ξ is coorientable at Σ; it follows that $\mathcal{F}_\xi$ actually keeps track of the isomorphism class of $TM/\xi|_\Sigma$. Furthermore, we only used the orientability of Σ in a mild way: since M is orientable it follows that $TM_\Sigma \cong T\Sigma \oplus \mathbb{R}$ if and only if Σ is orientable. If Σ is non-orientable, $TM_\Sigma \cong T\Sigma \oplus \det T\Sigma$ (where $\det T\Sigma$ is the line bundle $T\Sigma \wedge T\Sigma$). Other than this small change, the proof above can be repeated verbatim for non-orientable Σ.

In some ways, this theorem ends the story about surfaces in contact 3-manifolds: there exists a contactomorphism $\varphi : \nu\Sigma_1 \to \nu\Sigma_2$ so that $\varphi(\Sigma_1) = \Sigma_2$ if and only if there is a diffeomorphism $\psi : \Sigma_1 \to \Sigma_2$ preserving characteristic foliations. However in some ways, the condition $\varphi(\Sigma_1) = \varphi(\Sigma_2)$ is very strong. For example, it may be that Σ_1 is embedded very non-generically with respect to the contact structure. Then it may be that $(\Sigma_1, \mathcal{F}_{\xi_1})$ is non-diffeomorphic to $(\Sigma_2, \mathcal{F}_{\xi_2})$, but after a C^∞ small perturbation of Σ_1 the foliations become diffeomorphic. It follows then that a neighborhood of Σ_2 admits a contact embedding into any neighborhood of Σ_1.

7.2 Generic singularities of characteristic foliations

We study singularities of the characteristic foliation of a generically embedded surface in a contact manifold. Because our considerations are local, we will assume that Σ is oriented and $\xi = \ker\alpha$ is cooriented. Proposition 7.2 then implies that $\mathcal{F}_\xi$ is determined by $\beta := \alpha|_\Sigma \in \Omega^1\Sigma$ up to multiplication by a positive function.

Choose a volume form σ on Σ matching the orientation, and define $X \in \Gamma(T\Sigma)$ by $X \lrcorner\, \sigma = \beta$. Since both β and σ are well defined up to scaling with a positive function (and since $\lrcorner$ is C^∞ linear), $\mathcal{F}_\xi$ determines X up to multiplication by a positive function. Notice that $X = 0$ exactly where $\mathcal{F}_\xi$ is singular, and on the non-singular set $\operatorname{span}(X) = \mathcal{F}_\xi$.

Therefore, singularities of characteristic foliations correspond to singularities of vector fields (or more precisely, since X is only defined up to rescaling, we are looking at singularities of integral curves of vector fields). The most basic properties of such a singularity are defined by the linearization of the vector field.

Let $X = f\partial_x + g\partial_y$ be a vector field on $\mathbb{R}^2$ that vanishes at the origin. The linearization of X at zero is the matrix

$$A = \begin{pmatrix} \frac{\partial f}{\partial x}(0,0) & \frac{\partial f}{\partial y}(0,0) \\ \frac{\partial g}{\partial x}(0,0) & \frac{\partial g}{\partial y}(0,0) \end{pmatrix}.$$

Let λ_1, λ_2 be the eigenvalues of A. Since a coordinate change of $\mathbb{R}^2$ has the effect of conjugating A, they only depend on X. We say that the singular point is *non-degenerate* if $\det A = \lambda_1\lambda_2 \neq 0$ (which implies in particular that the singularity is isolated). We say that $(0,0)$ is an *elliptic* singular point if $\lambda_1\lambda_2 > 0$, and *hyperbolic* if $\lambda_1\lambda_2 < 0$. Notice that, if $Y = hX$ for a nonvanishing function h on $\mathbb{R}^2$, then $A_Y = hA_X$, so the elliptic/hyperbolic distinction is actually an invariant of $\mathcal{F}$ (even after switching orientation).

This elliptic/hyperbolic dichotomy exists for arbitrary singular foliations, but in the contact case we have additional information. Since $\beta = \alpha|_\Sigma$ and $\alpha \wedge d\alpha \neq 0$, it follows that whenever $\beta = 0$, $d\beta \neq 0$. In the case where Σ and ξ are oriented, the sign of $d\beta$ tells us whether $T\Sigma = \xi$ with matching or opposite orientations. Since $d(X \lrcorner \sigma) = d\beta$, we can let $\sigma = dx \wedge dy$ on $\mathbb{R}^2$, and calculate that $d\beta = d(fdy - gdx) = (\frac{\partial f}{\partial x} + \frac{\partial g}{\partial y})dx \wedge dy = \mathrm{tr}(A)\sigma$. Thus in the case of singularities of characteristic foliations $\lambda_1 + \lambda_2 \neq 0$, and the sign tells us the sign of the tangency $T\Sigma = \xi$ (this is true even at degenerate singularities).

If $\Sigma \subseteq (M, \ker\alpha)$, we would like to know that a C^∞ generic perturbation of Σ causes the singularities of $\mathcal{F}_\xi$ to be non-degenerate. Of course a C^∞ generic vector field will have non-degenerate singularities, and since non-degeneracy is invariant under functional scaling of X, it follows that a C^∞ generic singular foliation will have non-degenerate singularities. Therefore a generic perturbation of ξ will create non-degenerate singularities. But by Gray stability this perturbation of ξ is realized by a smooth isotopy of Σ. Since a generic isotopy is not a contactomorphism (the set of contact vector fields has infinite codimension in the space of all vector fields) it follows that $\mathcal{F}_\xi$ will have non-degenerate singularities after C^∞ perturbing Σ in a generic way.

Thus, assuming that $\Sigma \subseteq (M, \xi)$ is embedded generically, all singularities will be isolated and fall into either the class of elliptic or hyperbolic singularities. Assuming that Σ is oriented and ξ is cooriented, each singularity can also be regarded as positive or negative. Let e_+, e_-, h_+, and h_- denote the total number of each type.

Proposition 7.6. *Let (M, ξ) be a cooriented contact manifold, and let $\Sigma \subseteq M$ be a closed oriented surface, generically immersed with respect to ξ. Then*

$$\chi(\Sigma) = (e_+ - h_+) + (e_- - h_-)$$

$$e(\xi|_\Sigma) = (e_+ - h_+) - (e_- - h_-).$$

Proof: Let X be as above: $X \lrcorner \sigma = \alpha|_\Sigma$ for some contact form α for ξ and volume form σ for Σ. Then X is a vector field in $T\Sigma$, and therefore the Poincaré-Hopf Theorem tells us that $\chi(\Sigma) = e(T\Sigma) = e_{\mathrm{tot}} - h_{\mathrm{tot}}$. For the second equation, notice that X is also a vector field in $\xi|_\Sigma$, so we can also apply the Poincaré-Hopf Theorem here. However at the negative singularities, the orientation of ξ is opposite to the orientation of $T\Sigma$, and therefore the contribution of those zeros will count oppositely. □

Remark 7.7. Another way to see the the elliptic/hyperbolic distinction is as follows. Suppose we are in the situation of Example 5.9, where $M = W \cap S^{2N-1}$

for some W of complex dimension 2 holomorphically embedded in $\mathbb{C}^N$. As we noted there, the 2-plane field $\xi = TM \cap \sqrt{-1}TM$ is a contact structure (equal to $\ker \lambda_{\mathrm{std}}|_M$).

Now if Σ is a surface in M, we immediately see that $T\Sigma_q = \xi_q$ at a point q if and only if $T\Sigma_q = \sqrt{-1}T\Sigma_q$. That is, $\mathcal{F}_\xi$ will be singular exactly where $T\Sigma$ is a complex (or anti-complex!) line.

Let $Gr_2^+(TW)$ be the bundle of oriented 2-dimensional subspaces in TW. (Since $TM \cong M \times \mathbb{R}^3$ we see that $TW|_M \cong M \times \mathbb{R}^4$, so in fact $Gr_2^+(TW)|_M \cong M \times Gr_{4,2}^+$, but we do not use this.) $T\Sigma$ therefore defines a section $T : \Sigma \to Gr_2^+(TW)|_\Sigma$. Sitting inside $Gr_2^+(TW)$ is the bundle $\mathbb{P}_\mathbb{C}(TW)$, the bundle of all complex lines in TW. Also there is a disjoint diffeomorphic subbundle $\mathbb{P}_{\overline{\mathbb{C}}}(TW) \subseteq Gr_2^+(TW)$, the bundle of anti-complex lines.

On can check that the fiber $Gr_{4,2}^+$ is diffeomorphic to $S^2 \times S^2$, but for our purposes if suffices to know that this manifold is 4-dimensional. Of course the fiber of $\mathbb{P}_\mathbb{C}(TW)$ – $\mathbb{CP}^1$ – is 2-dimensional, therefore in the generic case where $T(\Sigma) \pitchfork \mathbb{P}_\mathbb{C}(TW)$, they will intersect it in some finite number of points. The positive intersections are the positive elliptic singularities, and the negative intersections are the positive hyperbolic singularities. Similarly, the negative elliptic/hyperbolic singularities of $\mathcal{F}_\xi$ will correspond to the positive/negative intersections of $T(\Sigma)$ with $\mathbb{P}_{\overline{\mathbb{C}}}(TW)$.

The interesting part of this observation is that the embedding $\Sigma \subseteq W$ determines this information without reference to M, ξ or λ, it only uses the complex structure on W. This is perhaps surprising, because the foliation $\mathcal{F}_\xi$ itself cannot be determined without reference to more structure. Also note that this trick can be done for any contact manifold (M, ξ), by putting a cylindrical almost contact structure on the positive symplectization.

Remark 7.8. A singularity is C^1 diffeomorphic to its linearization, blah blah to be filled in later.

8 Curves in contact 3-manifolds

Having outlined the basics of surfaces in contact 3-manifolds, we now turn our attention to curves. There are two basic types of curves we can study: transverse and Legendrian.

8.1 Transverse knots

Definition 8.1. Let (M, ξ) be a contact manifold. A curve $\gamma \subseteq M$ is called transverse if $\gamma \pitchfork \xi$ everywhere.

Notice that ξ is cooriented near any transverse curve γ, since $T\gamma \cong TM/\xi|_\gamma$. This also says that a coorientation of ξ determines an orientation of γ. Unless stated otherwise we will assume that these orientations are chosen to match, and reserve the term *negatively transverse curve* for a transverse curve whose orientation is opposite of a given coorientation of ξ.

We describe a local neighborhood theorem of transverse curves.

Proposition 8.2. *Let γ be a connected transverse curve in (M, ξ). then there is an open neighborhood $U \supseteq \gamma$ and a contact embedding $\varphi : U \to (S^1 \times \mathbb{R}^2, \ker(dz - ydx + xdy))$, where $(x, y) \in \mathbb{R}^2$ and $z \in S^1 = \mathbb{R}/\mathbb{Z}$. This contact embedding satisfies $\varphi(\gamma) = \{x = y = 0\}$.*

Proof: By now this should be an easy exercise for the reader: find a diffeomorphism which preserves the contact structure at the curve γ, and then use Gray stability to construct a flow, giving a contactomorphism. The only subtlety is noting that ξ must be an orientable bundle (why?), and that any orientable 2-plane bundle over a circle is trivial. □

Proposition/Definition 8.3. Let $\gamma \subseteq (M, \xi)$ be a transverse curve in a coorientable contact manifold, and let F be a compact oriented surface with $\partial F = \gamma$. Since $\partial F \neq \varnothing$ it follows that $\xi|_F$ is a trivial 2-plane bundle, therefore we can find a non-zero vector field $X \in \Gamma(\xi|_F)$. Let γ_X be a small perturbation of γ in the direction of X. Then the linking number $\ell k(\gamma, \gamma_X)$ does not depend on the choice of vector field X. We call this number the *self-linking number* of γ, and denote it by $\ell(\gamma, F)$.

The proof that $\ell(\gamma, F)$ does not depend on X follows from

Proposition 8.4. *Let $\gamma \subseteq (M, \xi)$ be a transverse curve in a cooriented contact manifold, and let F a a compact oriented surface with $\partial F = \gamma$, which is generically embedded with respect to ξ. Then $\ell(\gamma, F) = -(e_+ - h_+) + (e_- - h_-)$.*

Proof: : Choose a collar neighborhood F_0 of γ in F so that $TF_0 \pitchfork \xi$ on the neighborhood, we identify this neighborhood with $F_0 = \gamma \times [0, \epsilon)$, so that the fibers $\{\text{point}\} \times [0, \epsilon)$ are the leaves of $\mathcal{F}_\xi$. Since ξ_F is trivial, we choose a trivialization $\xi_{F_0} \cong F_0 \times \mathbb{R}^2$ so that $\mathcal{F}_\xi = \{0\} \times \mathbb{R} \subseteq \mathbb{R}^2$.

Let $X \in \Gamma(\xi_{\mathcal{F}})$ be a non-vanishing vector field, with the above identification we identify X as a map $X : \gamma \times [0, \epsilon) \to \mathbb{R}^2$. By homotopy we can of course assume X only depends on the coordinate in γ. The linking number $\ell k(\gamma, \gamma_X)$ can be defined as the algebraic count of intersections of γ_X and F, therefore $\ell k(\gamma, \gamma_X)$ is equal to the algebraic intersection count of $X(\gamma)$ with $\{0\} \times [0, \infty)$ in $\mathbb{R}^2$. Equivalently, we can count the intersections of $X(\gamma)$ with $\{0\} \times (-\infty, 0]$.

Choose a bump function $f : [0, \epsilon) \to [0, 1]$ which is 0 near 0 and 1 near ϵ. Let $(q, t) \in \gamma \times [0, \epsilon)$, we define the vector field $\widetilde{X} : \gamma \times [0, \epsilon)$ by $\widetilde{X}(q, t) = (1 - f(t))(0, 1) + f(t)X(q)$. Since $\widetilde{X} = X$ outside of a compact subset, we extend $\widetilde{X}$ to all of F as being equal to X outside of F_0. Notice $\text{span}(\widetilde{X}_\gamma) = \text{span}(0, 1) = \mathcal{F}_\xi$. Notice also that for a fixed $q \in \gamma$, there is a $t \in [0, \epsilon)$ so that $\widetilde{X}(q, t) = 0$ if and only if $X(q) \in \{0\} \times (-\infty, 0]$. Therefore, the algebraic count of the zeroes of $\widetilde{X}$ is equal to $\pm \ell k(\gamma, \gamma_X)$.

Now let Y be a vector field defining $\mathcal{F}_\xi$, since we are free to scale Y we choose $Y = (0, 1)$ on F_0. Then $\widetilde{X}$ and Y are both vector fields in ξ_F with non-degenerate zeros, and they are equal outside of a compact set. It follows that $(e_+ - h_+) - (e_- - h_-) = \pm \ell k(\gamma, \gamma_X)$.

To calculate the sign ambiguity, note that the Poincaré-Hopf theorem states that the algebraic count of zeroes of $\widetilde{X}$ and Y are equal, with sign. The sign

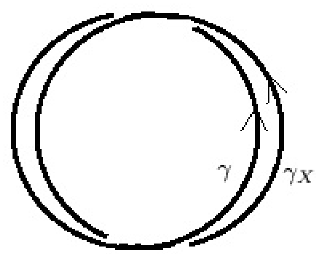

Fig. 1: Calculation of $\ell(\gamma, F)$ for a simple transverse curve in $\mathbb{R}^3_{\mathrm{std}}$.

ambiguity comes from the first part of the proof, when we compared the signed count of zeroes of $\widetilde{X}$ with $\ell k(\gamma, \gamma_X)$. So this sign ambiguity happens completely in a local neighborhood of γ, independent of F. Thus we can calculate this sign for all transverse curves simply by calculating it in one example.

To this end consider $(\mathbb{R}^3, \ker(dz + r^2 d\theta))$, let $\gamma = \{z = 0, r = C\}$, and $F = \{z = 0, r \leqslant C\}$. Then with the coorientation given by $dz + r^2 d\theta$, γ is oriented counterclockwise, and thus F is cooriented by ∂_z. On $\mathcal{F}_\xi$ has a unique singularity at $r = 0$, and it is a positive elliptic. The projection $\xi \to TF$ is an isomorphism, so let X be the vector field in ξ which projects to ∂_x. Then γ_X looks as in Figure 1, and we see that $\ell(\gamma, F) = -1$. □

Remark 8.5. Suppose we switch the orientation of the curve γ. This will then switch the orientation of F, but we should also switch the coorientation of ξ to match γ. Then the numbers e_+, h_+, e_-, h_- do not change, since we have switched the orientations of both TF and ξ. Therefore, *the invariant $\ell(\gamma, F)$ does not depend on orientation choices, as long as the orientations are assumed to be compatible.*

Suppose we are given two surfaces F_1, F_2 both of which have γ as a boundary. The surface framing of a nulhomologous smooth knot is well defined in any 3-manifold (this follows from Poincare duality), therefore after smoothing corners we can glue the surfaces together to get a smooth (immersed) surface $F = F_1 \cup_\gamma -F_2$. Since γ is transverse this corner smoothing does not induce any new singularities of $\mathcal{F}_\xi$. Therefore, Propositions 8.4 and 7.6 together imply

$$\ell(\gamma, F_1) - \ell(\gamma, F_2) = e(\xi)[-F_1 \cup_\gamma F_2].$$

In particular if $e(\xi) = 0 \in H^2(M)$, then $\ell(\gamma)$ is a well defined integer that does not depend on the choice of F.

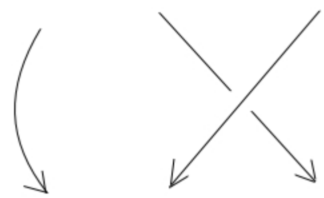

Fig. 2: The two local pictures which cannot exist as the (x, z)-projection of a transverse curve.

Remark 8.6. Just as in Proposition 7.6, it follows immediately from the Poincaré-Hopf theorem that $\chi(F) = (e_+ - h_+) + (e_- - h_-)$ for a surface F with transverse boundary. When we calculate modulo 2, we have that $\ell(\partial F, F) \equiv \chi(F) \equiv \dim H_0(\partial F) \mod 2$. In particular if ∂F is connected then the self-linking number is always odd.

We now give a proposition that describes how to draw pictures of transverse curves in $(\mathbb{R}^3, \xi_{\text{std}} = \ker(dz - ydx))$.

Proposition 8.7. *Let D be an oriented knot diagram in $\mathbb{R}^2$ (a self-transverse immersed curve with self-intersections assigned as over or under crossings). Assume that D is never oriented downward at a vertical tangency, and that it has no crossings as in Figure 2. Then D is the (x, z) projection of a transverse curve in $\mathbb{R}^3_{std}$, and that transverse curve is unique up to isotopy.*

Proof: Let $\gamma : S^1 \to \mathbb{R}^3_{\text{std}}$ be a parametrized curve, we write it in components $\gamma = (\gamma_x, \gamma_y, \gamma_z)$. Then γ is a transverse curve iff $\dot{\gamma}_z - \gamma_y \dot{\gamma}_x > 0$ at every point. A diagram D defines the functions γ_x and γ_z, therefore the proposition is interpreted as existence and uniqueness of suitable functions γ_y.

First let us discuss existence. If $\dot{\gamma}_x = 0$ then we must have $\dot{\gamma}_z > 0$ in order to satisfy the transverse condition, thus there can be no downward vertical tangencies. Elsewhere the condition becomes

$$\gamma_y < \frac{\dot{\gamma}_z}{\dot{\gamma}_x} \text{ if } \dot{\gamma}_x > 0$$

$$\gamma_y > \frac{\dot{\gamma}_z}{\dot{\gamma}_x} \text{ if } \dot{\gamma}_x < 0.$$

Of course this can be solved away from the crossings of D. But the labeling of under/over crossings of D gives us additional constraints, of the form $\gamma_y(t_1) > \gamma_y(t_2)$.

If the sign of $\dot{\gamma}_x$ is equal for the two branches of a crossing then it is easy to solve these constraints. If $\dot{\gamma}_x$ takes opposite signs and the crossing is oriented upward then we have

$$\gamma_y(t_1) < \frac{\dot{\gamma}_z}{\dot{\gamma}_x}(t_1) > \frac{\dot{\gamma}_z}{\dot{\gamma}_x}(t_2) < \gamma_y(t_2)$$

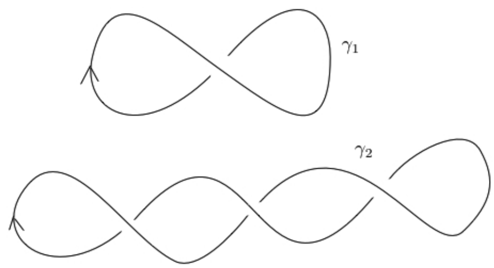

Fig. 3: Because all crossings in both pictures are negative, we see that $\ell(\gamma_1) = -1$ and $\ell(\gamma_2) = -3$.

and this equation we can solve with either ordering on the set $\{\gamma_y(t_1), \gamma_y(t_2)\}$. But if $\dot\gamma_x$ takes opposite signs and the crossing is oriented downwards then we must solve

$$\gamma_y(t_1) < \frac{\dot\gamma_z}{\dot\gamma_x}(t_1) < \frac{\dot\gamma_z}{\dot\gamma_x}(t_2) < \gamma_y(t_2)$$

and therefore the sign of the crossing is constrained. To see that the crossing in Figure 2 is the one which is not allowed notice that the y axis points *into* the page, since $dx \wedge dy \wedge dz$ is the canonical orientation on $\mathbb{R}^3_{\text{std}}$.

Finally to show uniqueness, note that if γ_y^0 and γ_y^1 are two transverse liftings of the diagram D, then $\gamma_y^s = (1-s)\gamma_y^0 + s\gamma_y^1$ defines an isotopy through transverse curves. □

Finally, we note how to calculate the self-linking number of a transverse curve in $\mathbb{R}^3_{\text{std}}$. But this is easy: the vector field ∂_y is nonzero and contained in ξ at every point, and therefore the self-linking number is equal to the blackboard framing of the (x, z)-projection. This is equal to the writhe of the diagram: the total number of positive crossings minus negative crossings.

Example 8.8. Figure 3 gives pictures of two transverse knots which are smoothly unknotted, and calculates their self-linking number. By generalizing γ_2 to a knot with more twists, it's easy to construct transverse unknots with any negative odd self-linking number. Thus there are infinitely many transverse unknots which are mutually distinct as transverse curves.

8.2 Legendrian knots

We discuss how to draw pictures of Legendrian curves $\Lambda \subseteq \mathbb{R}^3_{\text{std}}$. Just as in the case of transverse curves, we will work in the (x, z) projection, which we denote by $\pi : \mathbb{R}^3_{\text{std}} \to \mathbb{R}^2$. Writing $\Lambda(t) = (\Lambda_x(t), \Lambda_y(t), \Lambda_z(t))$, the Legendrian condition is given by $\dot\Lambda_z = \Lambda_y \dot\Lambda_x$. In particular $\pi(\Lambda)$ cannot have vertical tangencies anywhere, and whenever $\pi(\Lambda)$ is smooth we can recover the y component of Λ by $\Lambda_y = \frac{\dot\Lambda_z}{\dot\Lambda_x}$.

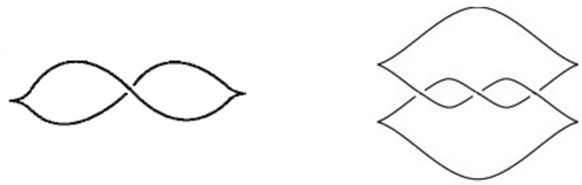

Fig. 4: Pictures of Legendrian curves in $\mathbb{R}^3_{\text{std}}$. The first curve is smoothly unknotted, the second is smoothly isotopic to the trefoil knot.

Because $\pi(\Lambda)$ can have no vertical tangencies it may seem impossible to have a closed Legendrian curve in $\mathbb{R}^3_{\text{std}}$. But in fact they exist in abundance, the issue is that $\pi(\Lambda)$ will not be a smooth curve, even when Λ is generically embedded. Consider the curve $\Lambda(t) = (t^2, \frac{3}{2}t, t^3)$. Then Λ is a smoothly embedded Legendrian curve, but $\pi(\Lambda)$ is singular at $t = 0$, having a *cusp* singularity (alternatively called a semi-cubic singularity). One can check that this singularity persists after any C^∞ perturbation of Λ. In fact this is the unique generic singularity of (x, z) projections of Legendrian curves (together with transverse self-intersections).

Because of the Legendrian condition, the over/under-crossings are uniquely determined by the projection $\pi(\Lambda)$: the branch with the smaller slope has smaller y-coordinate, and therefore this branch is the overcrossing in the diagram. See Figure 4 for some examples.

The next theorem says that the space of Legendrian curves is C^0 dense is the space of all smooth curves.

Theorem 8.9. *Let (M, ξ) be a contact 3-manifold, and let $\gamma : S^1 \to M$ be a smooth embedded curve. Then there exists a Legendrian curve $\Lambda : S^1 \to M$ which is arbitrarily C^0 close to γ, and smoothly isotopic to it. Furthermore, if γ is already Legendrian on a closed subset $A \subseteq S^1$ then we can arrange that $\Lambda = \gamma$ on A.*

We first prove a local lemma.

Lemma 8.10. *Let $\gamma : \mathbb{R} \to \mathbb{R}^3_{std}$ be a properly embedded curve which is Legendrian outside of a compact subset $K \subseteq \mathbb{R}^3_{std}$. Then there is a curve $\Lambda : \mathbb{R} \to \mathbb{R}^3_{std}$ which is Legendrian everywhere, C^0 close to γ and isotopic to γ via an isotopy supported in K.*

Proof: Let $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ be the projection onto the (x, z) coordinates. By C^∞ perturbation of $\gamma \cap K$ we may assume that $\pi(\gamma \cap K)$ is immersed. The y coordinates of γ define a line field η in $\pi^* T\mathbb{R}^2|_\gamma$, defined by $\eta = \ker(dz - ydx)$. Since γ is an embedded curve, it follows that at all double points $q, q' \in \gamma$, $\pi(q) = \pi(q')$, we have that $\eta_q \neq \eta_{q'}$ (unless $q = q'$). Let $V = \gamma \times (\epsilon, \epsilon)$ be the abstract normal bundle of $\pi(\gamma)$, that is $V = \pi^* T\mathbb{R}^2|_\gamma / T\gamma$, and let $\psi : V \to \mathbb{R}^2$

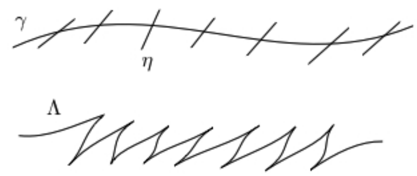

Fig. 5: The proof of Lemma 8.10.

be the immersion of the normal bundle, so that $\psi(\gamma \times \{0\}) = \pi(\gamma)$ and every point in $\psi(V)$ is within ϵ distance of $\pi(\gamma)$. We extend η to V by translation. Notice that, after possibly choosing a smaller ϵ, $q \mapsto (\pi(q), \eta_q)$ is an injective map from V to $\mathbb{R}^2 \times \mathbb{RP}^1$.

Notice that γ is Legendrian at a point $q \in \gamma$ if and only if $T\pi(\gamma)_q = \eta_q$. In particular, a Legendrian curve $\Lambda \subseteq \mathbb{R}^3_{\text{std}}$ is C^0 close to γ if and only if $\pi(\Lambda)$ is C^0 close to $\pi(\gamma)$ and $T\pi(\Lambda)$ is C^0 close to η. Instead of constructing $\Lambda \subseteq \mathbb{R}^3_{\text{std}}$ directly, we will construct its projection $\pi(\Lambda)$, which then has a unique lift to $\mathbb{R}^3$ as a Legendrian.

The basic idea of the proof is contained in Figure 5. It is impossible to approximate γ with a graphical Legendrian $\{z = f(x), y = f'(x)\}$, but by choosing a non-graphical Legendrian with cusp singularities we can get around this. Let $L = \{q \in \gamma; |\,\text{slope}(T\pi(\gamma)_q) - \text{slope}(\eta_q)| \leqslant \epsilon\}$ (if $T\pi(\gamma)_q$ is vertical then $q \notin L$). Over L, we simply define $\pi(\Lambda) = \pi(\gamma)$. Then Λ is C^0 close to γ, and $\Lambda = \gamma$ outside of K. $\gamma \setminus L$ is then a union of open intervals, we denote one such interval by (T_0, T_1). It suffices to show that we can approximate γ by Λ on this interval.

Let $V_0 \subseteq V$ be the portion of the normal bundle over (T_0, T_1). Because η is transverse to γ on V_0, we can integrate η to give coordinates (q, r) on V_0 where $q \in \gamma$ and ∂_r is tangent to η. We orient r so that ∂_r and ∂_q are ϵ-close to parallel on ∂V_0. Let $t_k = T_0 + \frac{k}{N}(T_1 - T_0)$, where N is a large integer we choose later. Let v_k be the intervals $v_k = \{q = t_k, |r| < \epsilon\}$, and let w_k be a smooth curve connecting the endpoints of v_k and v_{k+1}. See Figure 6.

Notice that the q/r slope of the curves w_k can be bounded above by $2\frac{(T_1 - T_0)}{N\epsilon}$, by rough estimation. Then define $\pi(\Lambda)$ to be the union of the curves $\{w_k\}$ (oriented as r), and $\{v_k\}$ (oriented as $-r$). Near the boundary, we cut off $\pi(\Lambda)$ however we like to match the given boundary conditions. Because ∂_q and ∂_r are nearly parallel near $\partial(T_0, T_1)$, $T\pi(\Lambda)$ will be ϵ close to η for any cutoff function we choose. The distortion between q/r slope and $z/x - \eta$ is bounded by a constant C, since it is bounded by ϵ near $\partial(T_0, T_1)$. Then, by choosing $N > 2C(T_1 - T_0)\epsilon^{-2}$, we get that $T\pi(\Lambda)$ is ϵ close to η everywhere. Because near the endpoints $T\pi(\Lambda)$ is ϵ-close to eta, this is also true near the endpoints for any increasing cutoff function we choose. □

Definition 8.11. Let $\Lambda \subseteq (M, \xi)$ be a nulhomologous Legendrian knot, so that

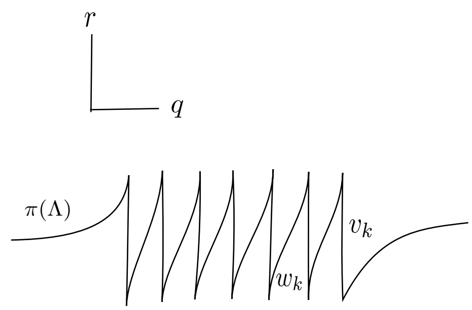

Fig. 6: The construction of $\pi(\Lambda)$ in q, r coordinates.

ξ is coorientable near Λ. Let X be a non-vanishing vector field along Λ which is transverse to ξ everywhere. The linking number of Λ with the X-pushoff of Λ is called the *Thurston-Bennequin number* of Λ, denoted $tb(\Lambda) = \ell k(\Lambda, \Lambda_X)$.

It is clear that $tb(\Lambda)$ does not depend on the choice of X: once we choose a coorientation of ξ, the set of X which are positively transverse to ξ make up a convex set. And furthermore, any X which is negatively transverse to ξ differs from one which is positively transverse to ξ by an 180 degree rotation.

Proposition/Definition 8.12. Let $\Lambda \subseteq (M, \xi)$ be a Legendrian knot, and F be an oriented surface so that $\partial F = \Lambda$. Assume that ξ is cooriented near F. Then $\xi|_F \cong F \times \mathbb{R}^2$. Then $q \mapsto T\Lambda_q$ is a map $S^1 \to \mathbb{R}^2 \setminus \{0\}$. The degree of this map is called the *rotation number* of Λ, denoted by $r(\Lambda, F)$. It does not depend on the choice of trivialization $\xi|_F \cong F \times \mathbb{R}^2$.

Proof: Given one trivialization $\xi|_F \cong F \times \mathbb{R}^2$, another (orientation preserving) trivialization is given by a map $f : F \to SO_2$. Since $[\Lambda] = 0 \in H_1(F)$, we have that $[f(\Lambda)] = 0 \in H_1(SO_2)$. But since Λ and SO_2 are both diffeomorphic to the circle, homotopy classes of maps $\Lambda \to SO_2$ are determined by H_1. Therefore any re-trivialization of ξ along F restricts to a re-trivialization of $\xi|_\Lambda$ which is homotopic to the original one. The result follows. □

Though the choice of trivialization does not affect $r(\Lambda)$, the choice of orientation on Λ and also the choice of coorientation of ξ do affect the sign of $r(\Lambda, F)$. Furthermore, the choice of F affects $r(\Lambda, F)$ as the euler class of ξ:

Proposition 8.13. *Let $\Lambda \subseteq (M, \xi)$ be a oriented Legendrian knot in a coorientable contact manifold. Suppose $\partial F_1 = \Lambda = \partial F_2$. Then $r(\Lambda, F_1) - r(\Lambda, F_2) = e(\xi)[F_1 \cup_\Lambda -F_2]$.*

Proof:

Definition 8.14. stabilization

Proposition 8.15. *Let $\Lambda \subseteq (M, \xi)$ be a nulhomologous Legendrian knot, where ξ is cooriented near Λ. Then $tb(s_+(\Lambda)) = tb(s_-(\Lambda)) = tb(\Lambda) - 1$. Suppose* *further that $\Lambda = \partial F$ is oriented and ξ is cooriented near F. Then $r(s_\pm(\Lambda), F) = r(\Lambda) \pm 1$.*

Proof:

Theorem 8.16 (Fuchs-Tabachnikov, 1995)**.** *Let $\Lambda_0, \Lambda_1 \subseteq (M, \xi)$ be Legendrian knots which are smoothly isotopic. Then for sufficiently large integers m_0, m_1, n_0, n_1, $s_+^{m_0} s_-^{n_0}(\Lambda_0)$ is Legendrian isotopic to $s_+^{m_1} s_-^{n_1}(\Lambda_1)$.*

Proof:

8.3 Transverse pushoffs of Legendrians

Proposition/Definition 8.17. Let $\Lambda \subseteq (M, \xi)$ be a Legendrian knot so that ξ is cooriented near Λ. Suppose Λ is equipped with an orientation. Then there exists a transverse knot γ which is C^1-close to Λ, and furthermore γ is uniquely defined up to isotopy.

Proof: Choose 1-Jet coordinates $\mathcal{J}^1(S^1)$ around $\Lambda = \{p = z = 0\}$. Any smooth curve γ which is C^1-close to Λ is graphical: $\gamma = \{z = z(q), p = p(q)\}$. Then the condition for γ to be transverse is $z'(q) > p(q)$. This set is non-empty (since $z = 0, p = -\epsilon$ works) and convex (hence connected). □

Corollary 8.18. *Let $c : S^1 \to (M, \xi)$ be a smooth embedded curve, so that ξ is coorientable near $c(S^1)$. Then there is a transverse curve γ which is C^0 close to c and smoothly isotopic to it.*

Notice that in order to define the transverse pushoff, we need both an orientation of Λ and a coorientation of ξ near Λ. If we switch the coorientation of ξ but fix the orientation of Λ, the transverse pushoff will be a *different* curve (for example given by $\{z = 0, p = +\epsilon\}$). In the literature this is often called the *negative transverse push-off* of Λ.

Proposition 8.19. *Let $\Lambda = \partial F \subseteq (M, \xi)$ be a Legendrian curve which is oriented, so that ξ is cooriented near F. Let γ be the transverse push-off of Λ. Then $\ell(\gamma) = tb(\Lambda) - r(\Lambda)$.*

Proof:

Corollary 8.20. *Let $\Lambda = \partial F \subseteq (M, \xi)$ be an oriented Legendrian knot with ξ cooriented near F. Then $tb(\Lambda) + r(\Lambda, F) \equiv 1 \mod 2$.*

Proof: Let γ be the transverse pushoff of Λ. Then $tb(\Lambda) + r(\Lambda, F) \equiv \ell(\gamma) \mod 2$. By Remark 8.6 the latter number is always odd. □

9 Lutz-Martinet and overtwistedness

Up to this point we have not done any global contact geometry: we have only studied contact structures near points, curves, and surfaces. The most fundamental global questions about contact 3-manifolds is existence and uniqueness:

fixing a closed oriented 3-manifold M, does it admit a contact structure? And given two contact structures on M, under what hypotheses can we conclude they are isotopic?

Two contact structures are isotopic if and only if they are homotopic through contact structures: any isotopy trivially induces a homotopy, and Gray's theorem is precisely the converse. The most basic invariant of a contact structure is the *homotopy class of plane field*: obviously if ξ_1 and ξ_2 are not even homotopic as 2-plane distributions then they cannot possibly be homotopic as contact structures. We therefore refine our previous question:

Question 9.1. *Let M be a closed oriented 3-manifold, and let η be a 2-plane distribution on M. Does there exist a contact structure ξ on M which is homotopic to η as a 2-plane distribution (so that the orientation induced by ξ matches the orientation of M)? If ξ_1 and ξ_2 are two such contact structures, under what conditions can we ensure that ξ_1 and ξ_2 are homotopic as contact structures?*

Note. We think of the homotopy class of 2-plane field as an essentially uninteresting invariant, because it is defined in the world of algebraic topology and therefore cannot detect any interesting geometric phenomena. For example, we can identify homotopy classes of cooriented 2-plane fields as homotopy classes of non-vanishing vector fields. Choosing a trivialization $TM \cong M \times \mathbb{R}^3$ (which always exists but is non-canonical, see the note at the beginning of Section 7) we can therefore identify homotopy classes of cooriented plane fields with homotopy classes of maps $M \to S^2$.

Of course the euler class of a cooriented contact structure is defined by the *isomorphism* class of 2-plane bundle, so it is an even weaker invariant than the homotopy class of 2-plane field. For example on $M = S^3$ the isomorphism class of a 2-plane bundle η is uniquely determined (since $\pi_2(SO_2) = \{1\}$), but homotopy classes of $\eta \subseteq TM$ are classified by $\pi_3(S^2) \cong \mathbb{Z}$. (This identification is not generally canonical since $TM \cong M \times \mathbb{R}^3$ is non-canonical, however in the special case of S^3 we can choose an identification by choosing a diffeomorphism $S^3 \cong SU_2$, since any Lie group is canonically parallelizable.)

The main goal of this section is to demonstrate the existence portion of the above question:

Theorem 9.2 (Lutz-Martinet, 1971)**.** *Let M^3 be a closed oriented 3-manifold, and let η be a 2-plane field on M. Then η is homotopic to a contact structure ξ, so that the orientation induced by ξ matches the given orientation on M.*

9.1 Overtwistedness and the Lutz tube

9.2 Proof of the Lutz-Martinet theorem

10 Hypersurfaces part II – Giroux flexibility

11 Giroux flexibility and tightness

11.1 Overtwistedness criterion and Bennequin's inequality

11.2 Manifolds with unique tight contact structures

11.3 Connect sum decompositions of tight contact manifolds

12 The classification of overtwisted contact structures

13 Bishop disk fillings